

Robust Online Motion Planning with Reachable Sets

Anirudha Majumdar, *Member, IEEE*,
Russ Tedrake, *Member, IEEE*

Abstract—In this paper we consider the problem of generating motion plans for a nonlinear dynamical system that are guaranteed to succeed despite uncertainty in the environment, parametric model uncertainty, disturbances, and/or errors in state estimation. Furthermore, we consider the case where these plans must be generated online, because constraints such as obstacles in the environment may not be known until they are perceived (with a noisy sensor) at runtime. Previous work on feedback motion planning for nonlinear systems was limited to offline planning due to the computational cost of safety verification. Here we augment the traditional *trajectory library* approach by designing locally stabilizing controllers for each nominal trajectory in the library and providing guarantees on the resulting closed loop systems. We leverage *sums-of-squares programming* to design these locally stabilizing controllers by explicitly attempting to minimize the size of the worst case reachable set of the closed-loop system subjected to bounded disturbances and uncertainty. The reachable sets associated with each trajectory in the library can be thought of as “funnels” that the system is guaranteed to remain within. The resulting *funnel library* is then used to *sequentially compose* motion plans at runtime while ensuring the safety of the robot. A major advantage of the work presented here is that by explicitly taking into account the effect of uncertainty, the robot can evaluate motion plans based on how vulnerable they are to disturbances. We demonstrate our method on a simulation of a plane flying through a two dimensional forest of polygonal trees with parametric uncertainty and disturbances in the form of a bounded “cross-wind”.

Note to Practitioners—We are motivated by the need for planning algorithms for robots that are able to deal with uncertainty in the form of unknown or unmodeled dynamics, state estimation errors and obstacle positions that are unknown until runtime. Existing approaches to this problem typically either fail to provide formal guarantees on the behavior of the system subjected to disturbances and uncertainty or are unable to deal with *a priori* unknown environments. Our approach is to compute a set of “motion primitives” for which we can provide formal guarantees on the behavior of the closed-loop system. In particular, for each motion primitive, we can provide bounds on where the system may end up given that it starts off in some set of initial conditions. One can visualize these sets as “funnels” that the system is guaranteed to remain within (there is one funnel associated with each motion primitive). At runtime, when the robot encounters a novel environment, it can combine these motion primitives in order to plan its way safely through the environment. We provide an example of how this approach may be used on a model of an unmanned aerial vehicle (UAV) flying through a forest with unknown (but bounded) speed while subjected to a cross-wind. There are two main challenges that practitioners must overcome when applying this approach. First, the computational tools used for computing “funnels” can often run into numerical difficulties. Second, hardware implementation of the approach requires one to have a reasonably accurate model of the dynamics of the system and the uncertainty and

disturbances it is subjected to. While not the focus here, examples of applying related algorithms to physical systems can be found in [19] and [25]. We are currently experimenting with this approach for different application domains including humanoid locomotion.

Index Terms—Motion planning, feedback motion planning, robust control, reachability analysis, sums-of-squares programming.


I. INTRODUCTION

THE ability to plan and execute dynamic motions under uncertainty is a critical skill that our robots must have in order to perform useful tasks in the real world. Whether the robot is an unmanned aerial vehicle (UAV) flying at high speeds through a cluttered environment in the presence of wind gusts, a legged robot traversing rough terrain, or a micro-air vehicle with noisy on-board sensing, the inability to take into account disturbances, model uncertainty and state uncertainty can have disastrous consequences.

Motion planning has been the subject of significant research in the last few decades and has enjoyed a large degree of success in recent years. Planning algorithms like the Rapidly-exploring Randomized Tree (RRT) [16], RRT* [15], and related trajectory library approaches [17] [8] [30] can handle large state-space dimensions and complex differential constraints. These algorithms have been successfully demonstrated on a wide variety of hardware platforms [28] [27]. However, a significant limitation is their inability to explicitly reason about uncertainty and feedback. Modeling errors, state uncertainty and disturbances can lead to failure if the system deviates from the planned nominal trajectories. This issue is sketched in Figure 1(a), where a UAV attempting to fly through a forest with a heavy cross-wind gets blown off its planned nominal trajectory and crashes into a tree.

Recently, planning algorithms which explicitly take into account feedback control have been proposed. The LQR-Trees algorithm [31] creates a tree of locally stabilizing controllers which can take any initial condition in some bounded region in state space to the desired goal. The approach leverages *sums-of-squares programming* (SOS) [23] for computing regions of finite time invariance for the locally stabilizing controllers. However, LQR-Trees lack the ability to handle scenarios in which the task and environment are unknown till runtime: the offline precomputation of the tree does not take into account potential runtime constraints like obstacles, and an online implementation of the algorithm is computationally infeasible.

In this paper, we present a partial solution to this problem by combining trajectory libraries, feedback control, and tools from Lyapunov theory and algorithmic algebra in order to perform robust motion planning in the face of uncertainty.



(a) A plane deviating from its nominal planned trajectory due to a heavy cross-wind.

(b) The “funnel” of possible trajectories.

Fig. 1. Not accounting for uncertainty while planning motions can lead to disastrous consequences.

In particular, in the offline computation stage, we first design a finite library of open loop trajectories. For each trajectory in this library, we use *sums-of-squares programming* (SOS) to design a controller that explicitly attempts to minimize the size of the worst case reachable set of the system given a description of the uncertainty in the dynamics and bounded external disturbances. This control design procedure yields an outer approximation of the reachable set, which can be visualized as a “funnel” around the trajectory, that the closed loop system is guaranteed to remain within. A cartoon of such a funnel is shown in Figure 1(b). Finally, we provide a way of *sequentially composing* these robust motion plans online in order to operate in a provably safe manner.

One of the most important advantages that our approach affords us is the ability to choose between the motion primitives in our library in a way that takes into account the dynamic effects of uncertainty. Imagine a UAV flying through a forest that has to choose between two motion primitives: a highly dynamic roll maneuver that avoids the trees in front of the UAV by a large margin or a maneuver that involves flying straight while avoiding the trees only by a small distance. An approach that neglects the effects of disturbances and uncertainty may prefer the former maneuver since it avoids the trees by a large margin and is therefore “safer”. However, a more careful consideration of the two maneuvers could lead to a different conclusion: the dynamic roll maneuver is far more susceptible to wind gusts and state uncertainty than the second one. Thus, it may actually be more robust to execute the second motion primitive. Further, it may be possible that neither maneuver is guaranteed to succeed and it is safer to abort the mission and simply transition to a hover mode. Our approach allows robots to make these critical decisions, which are essential if robots are to move out of labs and operate in real-world environments.

II. RELATED WORK

The motion planning aspect of our approach draws inspiration from the vast body of work that is focused on computing motion primitives in the form of trajectory libraries. For example, trajectory libraries have been used in diverse applications such as humanoid balance control [17], autonomous ground vehicle navigation [27], and grasping [2] [7]. The Maneuver Automaton [8] attempts to capture the formal properties of trajectory libraries as a hybrid automaton, thus providing a unifying theoretical framework. Maneuver Automata have also been used for realtime motion planning with static and dynamic obstacles [9]. Further theoretical investigations have focused on the offline generation of diverse but sparse trajectories that ensure the robot’s ability to perform the necessary tasks online in an efficient manner [11]. More recently, tools from sub-modular sequence optimization have been leveraged in the optimization of the sequence and content of trajectories evaluated online [7].

The body of literature that deals with planning under uncertainty is also relevant to the work presented here [4] [24]. While these approaches generate motion plans that explicitly reason about the effect of uncertainty and disturbances on the behavior of the system, distributions over states (“belief states”) are typically approximated as Gaussians for computational efficiency and the true belief state is not tracked. Thus, in general, one does not have robustness guarantees. The approach we take here is to assume that disturbances/uncertainty are *bounded* and provide explicit bounds on the reachable set to facilitate safe operation of the system.

Robust motion planning has also been a very active area of research in the robotics community. Early work focused on the purely kinematic problem of planning paths through configuration space with “tubes” of specified radii around them such that all paths in the tube remained collision-free [12]. Recent work has focused on reasoning more explicitly

about the manner in which disturbances and uncertainties influence the dynamics of the robot, and is closer in spirit to the work presented here. In particular, [26] approaches the problem through dynamic programming on a model with disturbances by making use of the Maneuver Automaton framework mentioned earlier. However, the work does not take into account obstacles in the environment and does not provide or make use of any explicit guarantees on allowed deviations from the planned trajectories in the Maneuver Automaton. Another approach that is closely related to ours is Model Predictive Control with Tubes [21]. The idea is to solve the optimal control problem online with guaranteed “tubes” that the trajectories stay in. However, the method is limited to linear systems and convex constraints.

In [10], the authors design motion primitives for making a quadrotor perform an autonomous backflip. Reachable sets for the primitives are computed via a Hamilton-Jacobi-Bellman differential game formulation. However, a *predetermined* controller is employed for the reachability analysis instead of designing a controller that seeks to minimize the size of the reachable set (it is possible in principle to do this, but inconvenient in practice). More importantly, while their approach handles unsafe sets that the system is not allowed to enter, it is assumed that these sets are specified *a priori*. In this paper, we are concerned with scenarios in which unsafe sets (such as obstacles) are not specified until runtime and must thus be reasoned about *online*.

The approach that is perhaps most closely related to our work is the recent work presented in [22]. The authors propose a randomized planning algorithm in the spirit of RRTs that explicitly reasons about disturbances and uncertainty. Specifications of input to output stability with respect to disturbances provide a parameterization of “tubes” (analogous to our “funnels”) that can be composed together to generate motion plans that are collision-free. The factors that distinguish the approach we present in this paper from the one proposed in [22] are our focus on the *realtime* aspect of the problem and use of sums-of-squares programming as a way of computing reachable sets. In [22], the focus is on generating safe motion plans when the obstacle positions are known *a priori*. Further, we provide a general technique for computing and explicitly minimizing the size of tubes.

A critical component of the work presented here is the computation of “funnels” for nonlinear systems via Lyapunov functions. This idea, along with the metaphor of a “funnel”, was introduced to the robotics community in [5], where funnels were *sequentially composed* in order to produce dynamic behaviors in a robot. In recent years, sums-of-squares programming has emerged as a way of checking the Lyapunov function conditions associated with each funnel [23]. The technique relies on the ability to check nonnegativity of multivariate polynomials by expressing them as a sum of squares of polynomials. This can be written as a semi-definite optimization program and is amenable to efficient computational algorithms such as interior point methods [23]. Assuming polynomial dynamics, one can check that a polynomial Lyapunov candidate, $V(x)$, satisfies $V(x) > 0$ and $\dot{V}(x) < 0$ in some region B_r . Importantly, the same idea can be

used for designing controllers along time-indexed trajectories of a system that attempt to maximize the size of the set of initial conditions that are driven to a goal set [19]. In this paper, we extend this approach to compute controllers that explicitly minimize the size of reachable sets around trajectories. Thus, we are guaranteed that if the system starts off in the set of given initial conditions, it will remain in the computed “funnel” even if the model of the dynamics is uncertain and the system is subjected to bounded disturbances and state uncertainty.

An alternative approach to computing outer approximations of reachable sets is the one presented in [13]. The method relies on computing regions of finite time invariance using locally valid “barrier functions”. Although the approach does not involve computing controllers that attempt to minimize the size of the reachable set, it is conceivable that the method could be extended to do so.

III. CONTRIBUTIONS

This paper makes two main contributions. First, we provide a way of designing controllers using sums-of-squares programming that explicitly seek to minimize the effect that disturbances and uncertainties have on the system by minimizing the size of the reachable set (“funnel”). These controllers and corresponding reachable set guarantees can be generated for time-varying polynomial systems subjected to a broad class of uncertainties (bounded uncertainty in parameters entering polynomially in the dynamics). This is an extension of results presented in [19], where the control design approach seeks to maximize the size of the set of initial conditions that are guaranteed to be driven to some predefined goal set. The present work extends this approach to handle disturbances/uncertainty and provide guarantees on reachable sets rather than the set of initial conditions that are driven to the goal set.

Second, we show how a library of such funnels can be precomputed offline and composed together at runtime in a receding horizon manner while ensuring that the resulting closed loop system is “safe” (i.e. avoids obstacles and switches between the planned sequence of funnels). This can be viewed as an extension of the LQR-Trees algorithm [31] for feedback motion planning, which was limited to offline planning due to the relatively large computational cost of computing the funnels. In contrast to LQR-Trees, our algorithm is suitable for real-time, online planning. We expect this framework to be useful in robotic tasks where the dynamics and perceptual system of the robot are difficult to model perfectly and for which the robot does not have access to the geometry of the environment until runtime.

IV. COMPUTING REACHABLE SETS

A considerable amount of research effort in the motion planning community has focused on the design of trajectory libraries (see Section II). Hence, here we assume that we are provided with a trajectory library consisting of a finite set of nominal feasible trajectories for the robot and concentrate our discussion on extending the techniques for the computation of controllers and associated regions of finite time invariance

presented in [19] to compute reachable sets when there is uncertainty in the dynamics and state. Let

$$\dot{x} = f(x(t), w(t)) + g(x(t), w(t))u(t)$$

be the control system under consideration. Here, $x(t) \in \mathbb{R}^n$ is the state, $u(t) \in \mathbb{R}^m$ is the control input and $w(t) \in \mathbb{R}^d$ is the disturbance/uncertainty term. We assume here that f and g are polynomials¹ in x and w . We further assume that $w(t)$ belongs to a bounded semialgebraic set $W = \{w \mid W_k(w) \geq 0, \forall k = 1, \dots, K\}$.

Let $x_0(t) : [0, T] \mapsto \mathbb{R}^n$ be a nominal trajectory in our library that we want the system to follow and $u_0(t) : [0, T] \mapsto \mathbb{R}^m$ be the corresponding nominal open-loop control input. Defining new coordinates $\bar{x} = x - x_0(t)$ and $\bar{u} = u - u_0(t)$, we can rewrite the dynamics in these variables as $\dot{\bar{x}} = \bar{x} - \dot{x}_0(t)$. Then, given a set of initial conditions $F(0)$, we seek to design a controller that attempts to minimize the “size” of the time-varying reachable sets $B(t)$ (we will formalize what we mean by “size” soon). For a given controller, the reachable set $B(t)$ is the set of states that the system may be driven to at time t by *some* disturbance (i.e. some choice of $w(t) \in W$), given that the initial condition lay in the set $F(0)$. In general, we will not be able to compute reachable sets exactly. Rather, we will compute *outer approximations* of the reachable sets and design controllers to minimize the “size” of the outer approximation. Checking the following *invariance* condition for all $t \in [0, T]$ is sufficient for establishing the sets $F(t)$ as outer approximations of the reachable sets $B(t)$:

$$\bar{x}(0) \in F(0) \implies \bar{x}(t) \in F(t), \forall w : [0, T] \rightarrow W. \quad (1)$$

Our task will be to design time-varying controllers that minimize the size of the “funnel” described by the sets $F(t)$. We describe the funnel as a time-varying sub-level set of a function $V(\bar{x}, t)$:

$$F(t) = \{\bar{x} \mid \bar{x} \in \mathbb{R}^n, V(\bar{x}, t) \leq \rho(t)\}.$$

This specification of the funnel allows us to use $\rho(t)$ as a natural surrogate for the “size” of the funnel at time t . We impose the following condition on $V(\bar{x}, t)$:

$$V(\bar{x}, t) = \rho(t) \implies \dot{V}(\bar{x}, t, w) < \dot{\rho}(t), \forall w(t) \in W \quad (2)$$

Letting $F(0) = \{\bar{x} \mid V(\bar{x}, 0) \leq \rho_0\}$, it is easy to see that this condition implies the invariance condition (1). Here, $\dot{V}(\bar{x}, t, w)$ is computed as:

$$\dot{V}(\bar{x}, t, w) = \frac{\partial V(\bar{x}, t)}{\partial \bar{x}} \dot{\bar{x}} + \frac{\partial V(\bar{x}, t)}{\partial t}.$$

In principle, we can parameterize our function $V(\bar{x}, t)$ as a polynomial in both t and x and check (2) $\forall t \in [0, T]$. However, as described in [32], this leads to expensive sums-of-squares programs. Instead, we can get large computational gains with little loss in accuracy by checking (2) at sample points in

time $t_i \in [0, T], i = 1 \dots N$. As discussed in [32], for a fixed $V(\bar{x}, t)$ and dynamics (and under mild conditions on both), increasing the density of the sample points eventually recovers (2) $\forall t \in [0, T]$. This allows us to *check* the answers we obtain from the sums-of-squares program below by sampling finely enough.

Thus, we parameterize $V(\bar{x}, t)$ and \bar{u} by polynomials $V(\bar{x}, t_i)$ and $\bar{u}(\bar{x}, t_i)$ respectively at each sample point in time. Using $\sum_{i=1}^N \rho(t_i)$ as the cost function, we can write the following sums-of-squares (SOS) program:

$$\begin{aligned} & \text{minimize} && \sum_{i=1}^N \rho(t_i) \\ & \rho(t_i), L(\bar{x}, t_i, w), V(\bar{x}, t_i), \bar{u}(\bar{x}, t_i), M_k(\bar{x}, t_i, w) \end{aligned} \quad (3)$$

subject to :

$$V(\bar{x}, t_i) \text{ is SOS}, \forall i = 1 \dots N \quad (4)$$

$$-\dot{V}(\bar{x}, t_i, w) + \dot{\rho}(t_i) + L(\bar{x}, t_i, w)(V(\bar{x}, t_i) - \rho(t_i)) \dots$$

$$- \sum_{k=1}^K M_k(\bar{x}, t_i, w) W_k(w) \text{ is SOS}, \forall i = 1 \dots N \quad (5)$$

$$M_k(\bar{x}, t_i, w) \text{ is SOS}, \forall i = 1 \dots N \quad (6)$$

$$\rho(t_i) \geq 0, \forall i = 2 \dots N \quad (7)$$

$$V(e, t_i) = V_{\text{guess}}(e, t_i), \forall i = 1 \dots N \quad (8)$$

Here, $L(\bar{x}, t_i, w)$ and $M_k(\bar{x}, t_i, w)$ are “multiplier” terms that help to enforce the invariance condition. It is easy to see that condition (5) is a sufficient condition for ensuring (2) at the sample points in time. This is because for all $w \in W$, we must have $\sum_k M_k(\bar{x}, t_i, w) W_k(w) \geq 0$, since we have $M_k(\bar{x}, t_i, w) \geq 0$ and $W_k(w) \geq 0$. Thus, when $V(\bar{x}, t_i) = \rho(t_i)$, condition (5) implies² that $\dot{V}(\bar{x}, t_i, w) < \dot{\rho}(t_i)$.

Condition (8) is a normalization constraint where e is the vector of all ones and $V_{\text{guess}}(\bar{x}, t)$ is the candidate for $V(\bar{x}, t)$ that is used to initialize the alternation scheme outlined below for solving the above optimization program. (If we do not impose a normalization constraint on $V(\bar{x}, t_i)$, $\sum_{i=1}^N \rho(t_i)$ can be made arbitrarily small simply by scaling the coefficients of $V(\bar{x}, t_i)$). We use a piecewise linear parameterization of $\rho(t)$ and can thus compute $\dot{\rho}(t_i) = \frac{\rho(t_{i+1}) - \rho(t_i)}{t_{i+1} - t_i}$. Similarly, we approximate $\frac{\partial V(\bar{x}, t_i)}{\partial t} \approx \frac{V(\bar{x}, t_{i+1}) - V(\bar{x}, t_i)}{t_{i+1} - t_i}$.

The above optimization program is not convex in general since it involves conditions that are bilinear in the decision variables. However, the conditions are linear in $L(\bar{x}, t_i, w)$, $\bar{u}(\bar{x}, t_i)$, $M_k(\bar{x}, t_i, w)$ for fixed $V(\bar{x}, t_i)$, $\rho(t_i)$, and are linear in $V(\bar{x}, t_i)$, $\rho(t_i)$, $M_k(\bar{x}, t_i, w)$ for fixed $L(\bar{x}, t_i, w)$, $\bar{u}(\bar{x}, t_i)$. Thus, in principle we could use a bilinear alternation scheme for solving this optimization program by alternating between the two sets of decision variables, $(L(\bar{x}, t_i, w), \bar{u}(\bar{x}, t_i), M_k(\bar{x}, t_i, w))$ and $(V(\bar{x}, t_i), \rho(t_i), M_k(\bar{x}, t_i, w))$ and repeat until convergence in the following two steps: (1) Fix $(V(\bar{x}, t_i), \rho(t_i))$ and search for $(L(\bar{x}, t_i, w), \bar{u}(\bar{x}, t_i), M_k(\bar{x}, t_i, w))$, and (2) Fix $(L(\bar{x}, t_i, w), \bar{u}(\bar{x}, t_i))$ and search for $(V(\bar{x}, t_i), \rho(t_i), M_k(\bar{x}, t_i, w))$. However, in the first step of this alternation, we cannot optimize the cost function $\sum_{i=1}^N \rho(t_i)$ since we have to fix $\rho(t_i)$ (we can optimize the cost function in the second step). We could

¹With the right change of coordinates, one can express the dynamics of most robotic systems as polynomials. For example, the dynamics of most rigid body systems can be transformed into polynomials by introducing new variables, s_i and c_i , for $\sin(\theta_i)$ and $\cos(\theta_i)$, and imposing the constraint that $s_i^2 + c_i^2 = 1$ (this equality constraint is easily imposed in the sums-of-squares programming framework). Another approach is to simply Taylor approximate the non-polynomial dynamics.

²SOS decompositions obtained from numerical solvers generically provide proofs of polynomial *positivity* as opposed to mere non-negativity (see the discussion in [1, p.41]). This is why we claim a strict inequality here.

simply make the first step a *feasibility* problem (instead of optimizing a cost function), but this prevents us from searching for a controller that explicitly seeks to minimize the desired cost function since in the second step of the alternation, we do not search for a controller. We get around this issue by introducing an additional step in the alternation, in which we fix $L(\bar{x}, t_i, w)$ and $V(\bar{x}, t_i)$ and search for $\bar{u}(\bar{x}, t_i)$, $\rho(t_i)$ and $M_k(\bar{x}, t_i, w)$, while minimizing $\sum_{i=1}^N \rho(t_i)$. The steps in the alternation are summarized in Algorithm 1.

Each iteration of the alternations in Algorithm 1 is guaranteed to obtain an objective $\sum_{i=1}^N \rho^*(t_i)$ that is at least as small as the previous one since a solution to the previous iteration is also valid for the current one.

Algorithm 1 Robust Controller Design

- 1: Initialize $V(\bar{x}, t_i)$ and $\rho(t_i)$, $\forall i = 1 \dots N$
 - 2: $\rho_{prev}(t_i) = 0$, $\forall i = 1 \dots N$.
 - 3: converged = false;
 - 4: **while** \neg converged **do**
 - 5: **STEP 1**: Solve feasibility problem by searching for $L(\bar{x}, t_i, w)$, $\bar{u}(\bar{x}, t_i)$, $M_k(\bar{x}, t_i, w)$, and fixing $V(\bar{x}, t_i)$, $\rho(t_i)$.
 - 6: **STEP 2**: Minimize $\sum_{i=1}^N \rho(t_i)$ by searching for $\bar{u}(\bar{x}, t_i)$, $\rho(t_i)$, $M_k(\bar{x}, t_i, w)$, and fixing $L(\bar{x}, t_i, w)$, $V(\bar{x}, t_i)$.
 - 7: **STEP 3**: Minimize $\sum_{i=1}^N \rho(t_i)$ by searching for $V(\bar{x}, t_i)$, $\rho(t_i)$, $M_k(\bar{x}, t_i, w)$, and fixing $L(\bar{x}, t_i, w)$, $\bar{u}(\bar{x}, t_i)$.
 - 8: **if** $\frac{\sum_{i=1}^N \rho(t_i) - \sum_{i=1}^N \rho_{prev}(t_i)}{\sum_{i=1}^N \rho_{prev}(t_i)} < \epsilon$ **then**
 - 9: converged = true;
 - 10: **end if**
 - 11: $\rho_{prev}(t_i) = \rho(t_i)$, $\forall i = 1 \dots N$.
 - 12: **end while**
-

Section IV-B discusses how to initialize $V(\bar{x}, t_i)$ and $\rho(t_i)$ for Algorithm 1.

A. Incorporating Actuator Limits

Our method allows us to incorporate actuator limits into the control design procedure. Although we examine the single-input case in this section, this framework is very easily extended to handle multiple inputs.

Let the control law $u(x)$ be mapped through the following control saturation function:

$$s(u(x)) = \begin{cases} u_{max} & \text{if } u(x) \geq u_{max} \\ u_{min} & \text{if } u(x) \leq u_{min} \\ u(x) & \text{o.w.} \end{cases}$$

where u_{max} and u_{min} are the maximum and minimum allowable inputs respectively. Then, a piecewise analysis of $\dot{V}(\bar{x}, t)$ can be used to check the Lyapunov conditions are satisfied even

when the control input saturates. Defining:

$$\begin{aligned} \dot{V}_{min}(\bar{x}, t, w) = & \frac{\partial V(\bar{x}, t)}{\partial \bar{x}}^T [f(\bar{x} + x_0(t), w) \dots \\ & + g(\bar{x} + x_0, w)u_{min} - \dot{x}_0(t)] + \frac{\partial V(\bar{x}, t)}{\partial t} \end{aligned} \quad (9)$$

$$\begin{aligned} \dot{V}_{max}(\bar{x}, t, w) = & \frac{\partial V(\bar{x}, t)}{\partial \bar{x}}^T [f(\bar{x} + x_0(t), w) \dots \\ & + g(\bar{x} + x_0, w)u_{max} - \dot{x}_0(t)] + \frac{\partial V(\bar{x}, t)}{\partial t} \end{aligned} \quad (10)$$

$$(11)$$

we must check the following conditions:

$$u(\bar{x}) \leq u_{min} \implies \dot{V}_{min}(\bar{x}, t) < \dot{\rho}(t) \quad (12)$$

$$u(\bar{x}) \geq u_{max} \implies \dot{V}_{max}(\bar{x}, t) < \dot{\rho}(t) \quad (13)$$

$$u_{min} \leq u(\bar{x}) \leq u_{max} \implies \dot{V}(\bar{x}, t) < \dot{\rho}(t) \quad (14)$$

Algorithm 1 can be modified to enforce these conditions with extra multipliers in a manner identical to the one presented in [19]. This modification is relatively straightforward and we do not present it here.

B. Initializing $V(\bar{x}, t_i)$ and $\rho(t_i)$

Obtaining an initial guess for $V(\bar{x}, t_i)$ and $\rho(t_i)$ is an important part of Algorithm 1. In [31], the authors use the Lyapunov function candidate associated with a time-varying LQR controller. The control law is obtained by solving a Riccati differential equation:

$$-\dot{S}(t) = Q - S(t)B(t)R^{-1}B(t)^T S(t) + S(t)A(t) + A(t)^T S(t)$$

with final value conditions $S(t) = S_f$. Here $A(t)$ and $B(t)$ describe the time-varying linearization of the dynamics about the nominal trajectory $x_0(t)$. Q and R are positive-definite cost-matrices. The function:

$$V_{guess}(\bar{x}, t) = (x - x_0(t))^T S(t) (x - x_0(t)) = \bar{x}^T S(t) \bar{x}$$

is our initial Lyapunov candidate. $V_{guess}(\bar{x}, t_N) = \bar{x}^T S(0) \bar{x}$, along with a choice of ρ_0 can be used to determine the initial condition set, $F(0)$ (ref. Section IV):

$$F(0) = \{\bar{x} \mid \bar{x} \in \mathbb{R}^n, \bar{x}^T S(0) \bar{x} \leq \rho_0\}.$$

We find that initializing $\rho(t_i)$ to an exponential function in time, $e^{\gamma t_i}$, works quite well in practice. We can tune γ to obtain feasible solutions. Intuitively, higher values of γ correspond to “larger” reachable sets and thus are more likely to be feasible.

V. FUNNEL LIBRARIES

The tools from Section IV can be used to create libraries of funnels offline. Given a trajectory library, \mathcal{T} , consisting of finitely many trajectories $x_i(t)$, we can compute robust controllers $u_i(x, t)$ and associated reachable sets (funnels) for each trajectory in \mathcal{T} . However, there is an important issue that needs to be addressed when designing libraries of funnels and has an analogy in the traditional trajectory library approach. In particular, trajectories in a traditional trajectory library need to be designed in a way that allows them to be sequenced together. More formally, let \mathcal{P} denote the projection operator

that projects a state, x , onto the subspace formed by the non-cyclic dimensions of the system (i.e. the dimensions with respect to which the dynamics of the system are *not* invariant). Then, for two trajectories $x_i(t)$ and $x_j(t)$ to be executed one after another, we must have

$$\mathcal{P}(x_i(T_i)) = \mathcal{P}(x_j(0)).$$

Note that the cyclic coordinates do not pose a problem since one can simply “shift” trajectories around in these dimensions. This issue is discussed thoroughly in [8] and is addressed by having a *trim trajectory* of the system that other trajectories (*maneuvers*) start from and end at (of course, one may also have more than one *trim trajectory*).

In the case of *funnel* libraries, however, it is neither necessary nor sufficient for the nominal trajectories to line up in the non-cyclic coordinates. It is the *interface* between funnels that is important. Let $x_i(t)$ and $x_j(t)$ be two nominal trajectories in our library and $F_i(t) = \{x \mid x \in \mathbb{R}^n, V_i(\bar{x}, t) \leq \rho_i(t)\}$ and $F_j(t) = \{x \mid x \in \mathbb{R}^n, V_j(\bar{x}, t) \leq \rho_j(t)\}$ be the corresponding funnels. Further, we write $x = [x_c, x_{nc}]$, where x_c represent the cyclic dimensions and x_{nc} the non-cyclic ones. We say that F_i is *sequentially composable* with F_j if

$$\begin{aligned} \mathcal{P}(F_i(T_i)) &\subset \mathcal{P}(F_j(0)) \\ \iff \forall x = [x_c, x_{nc}] \in F_i(T_i), \quad \exists x_{0,c} \text{ s.t. } [x_{0,c}, x_{nc}] &\in F_j(0). \end{aligned} \quad (15)$$

While (15) is a sufficient condition for two funnels to be executed one after another, the dependence of $x_{0,c}$ on x makes searching for $x_{0,c}$ a non-convex problem in general. Thus, we set $x_{0,c}$ to be the cyclic coordinates of $x_j(0)$, resulting in a stronger sufficient condition that can be checked easily via a sums-of-squares program:

$$\forall x = [x_c, x_{nc}] \in F_i(T_i), \quad [x_{0,c}, x_{nc}] \in F_j(0). \quad (16)$$

Intuitively, (16) corresponds to “shifting” the inlet of funnel F_j along the cyclic dimensions so it lines up with x_c . Note that not all pairs of funnels in the library will be sequentially composable in general. Thus, as we discuss in Section VI, we must be careful to ensure sequential composability when planning sequences of funnels online.

VI. ONLINE PLANNING WITH FUNNELS

Having computed libraries of funnels in the offline pre-computation stage, we can proceed to use these primitives to perform robust motion planning online. The robot’s task specification may be in terms of a goal region that must be reached (as in the case of a manipulator arm grasping an object), or in terms of a nominal direction the robot should move in while avoiding obstacles (as in the case of a UAV flying through a forest or a legged robot walking over rough terrain). For the sake of concreteness, we adopt the latter task specification although one can easily adapt the contents of this section to the former specification. We further assume that the robot is provided with polytopic regions in configuration space that obstacles are guaranteed to lie in and that the robot’s sensors only provide this information up to a finite (but receding) spatial horizon. Our task is to sequentially compose funnels from our library in a way that avoids obstacles while

moving forwards in the nominal direction. The finite horizon of the robot’s sensors along with the computational power at our disposal determines how long the sequence of planned funnels can be at any given time.

The most important computation that needs to be performed at runtime is to check whether a given funnel intersects an obstacle. For the important case in which our Lyapunov functions are quadratic in x , this computation is a Quadratic Program (QP) and can be solved very efficiently (as evidenced by the success of larger scale QP formulations used in Model Predictive Control [6]). We denote $\bar{x} = x - x_0(t)$ as before, where $x_0(t)$ is the nominal trajectory. Let a particular obstacle be defined by half-plane constraints $A_j x \geq 0$ for $j = 1, \dots, M$. Note that A_j will typically be sparse since it will contain zeros in places corresponding to non-configuration space variables (like velocities). Then, for $i = 1, \dots, N$, we solve the following QP:

$$\begin{aligned} \underset{\bar{x}}{\text{minimize}} \quad & V(\bar{x}, t_i) \\ \text{subject to} \quad & A_j x \geq 0, \forall j \end{aligned} \quad (17)$$

Denoting the solution of (17) for a given t_i as $V^*(\bar{x}^*, t_i)$, the funnel does not intersect the obstacle if and only if $V^*(\bar{x}^*, t_i) > \rho(t_i), \forall t_i$. Multiple obstacles are handled by simply solving (17) for each obstacle. An important point that should be noted is that we do not require the obstacle regions to be convex. It is only required that they are represented as unions of convex sets. This allows us to handle situations where multiple polytopic regions overlap to form a non-convex region.

For higher order polynomial Lyapunov functions, one could check the following sums-of-squares conditions for all t_i :

$$\begin{aligned} V(\bar{x}, t_i) - \rho(t_i) - \sum_j L_j(\bar{x}) A_j(\bar{x} + x_0(t_i)) &\text{ is SOS} \\ L_j(\bar{x}) &\text{ is SOS, } \forall j = 1 \dots M \end{aligned} \quad (18)$$

However, these provide only sufficient conditions for non-collision. Thus, if the conditions in (18) are met, one is guaranteed that there is no intersection with the obstacle. The converse is not true in general. Further, depending on the state-space dimension of the robot, this optimization problem may be computationally expensive to solve online. Hence, for tasks in which online execution speed is crucial, one may need to restrict oneself to quadratic Lyapunov functions.

Algorithm 2 provides a sketch of the online planning loop. At every control cycle, the robot updates its state in the world along with the obstacle positions. It then checks to see if the sequence of funnels it is currently executing may lead to a collision with an obstacle (which should only be the case if the sensors report new obstacles). If so, it replans a sequence of funnels that can be executed from its current state and are collision-free. The *ReplanFunnels*(x, \mathcal{O}) subroutine assumes that funnel sequences that are *sequentially composable* in the sense of Section V have been ordered by preference during the precomputation stage. For example, for a navigation task, sequences may be ordered by how much progress the robot makes in some nominal direction. *ReplanFunnels*(x, \mathcal{O}) goes through funnel sequences and checks two things. First, it

checks that its current state is contained in the first funnel in the sequence (after appropriately shifting the funnel in the cyclic dimensions). Second, it checks that the sequence leads to no collisions with obstacles. The algorithm returns the first sequence of funnels that satisfies both criteria. Finally, the online planing loop computes which funnel of the current plan it is in and applies the corresponding control input $u_i(x, t, \text{internal})$.

Of course, several variations on Algorithm 2 are possible. In general, the funnel primitives provide a discrete action space which can be searched by any heuristic planner - the primary considerations here are the additional constraint of sequential composability and the moderately more significant cost of collision checking. In practice, it also may not be necessary to consider re-planning at the frequency of the control loop. Instead, longer sections of the plan may be executed before re-planning. Also, instead of choosing the most “preferred” collision-free sequence of funnels, one natural cost function which could guide the search is the minimum over t_i of $\frac{V^*(\bar{x}^*, t_i)}{\rho(t_i)}$. As before, $V^*(\bar{x}^*, t_i)$ is the solution of the QP (17). Since the 1-sublevel set of $\frac{V(\bar{x}, t_i)}{\rho(t_i)}$ corresponds to the funnel, maximizing this is a reasonable choice for choosing sequences of funnels.

Algorithm 2 Online Planning

```

1: Initialize current planned funnel sequence,
    $\mathcal{F} = \{F_1, F_2, \dots, F_n\}$ 
2: for  $t = 0, \dots$  do
3:    $\mathcal{O} \leftarrow$  Obstacles in sensor horizon
4:    $x \leftarrow$  Current state of robot
5:   Collision  $\leftarrow$  Check if  $\mathcal{F}$  collides with  $\mathcal{O}$  by solving
     QPs (17)
6:   if Collision then
7:      $\mathcal{F} \leftarrow \text{ReplanFunnels}(x, \mathcal{O})$ 
8:   end if
9:    $F.\text{current} \leftarrow F_i \in \mathcal{F}$  such that  $x \in F_i$ 
10:   $t.\text{internal} \leftarrow$  Internal time of  $F.\text{current}$ 
11:  Apply control  $u_i(x, t.\text{internal})$ 
12: end for
```

In order to initialize and replan the sequence of funnels \mathcal{F} , it is required that the current state be contained inside the first funnel in the sequence. Assuming perfect state estimates are available, this is easily checked. However, if perfect state information is not available, one needs to ensure that all possible states the system could be in lie inside the funnel. Assuming that measurement errors are bounded, one can use robust state estimation to provide worst-case bounds on the state estimate. For example, [18] provides a way of doing robust state estimation for polynomial discrete time systems via sums-of-squares programming. Given an outer approximation of the set of states the system could be in, one can check that the entire set is contained inside the funnel.

VII. EXAMPLE

We demonstrate our approach on a model of an aircraft flying in two dimensions through a forest of polygonal trees.




Fig. 2. Visualization of the system showing the coordinate system, polygonal obstacles, and “cross-wind”.

A pictorial depiction of the model is provided in Figure 2. The aircraft is constrained to move at a fixed forward speed and can control the second derivative of its yaw angle. We introduce uncertainty into the model by assuming that the speed of the plane is uncertain and time-varying and that there is a time-varying “cross-wind” whose magnitude is instantaneously bounded. The full non-linear dynamics of the system are then given by:

$$\mathbf{x} = \begin{bmatrix} x \\ y \\ \psi \\ \dot{\psi} \end{bmatrix}, \quad \dot{\mathbf{x}} = \begin{bmatrix} -v(t) \cos \psi \\ v(t) \sin \psi \\ \ddot{\psi} \\ u \end{bmatrix} + \begin{bmatrix} w(t) \\ 0 \\ 0 \\ 0 \end{bmatrix}. \quad (19)$$

with the speed of the plane $v(t) \in [9.5, 10.5]$ m/s and cross-wind $w(t) \in [-0.3, 0.3]$ m/s. The control input is bounded in the range $[-350, 350]$.

The plane’s trajectory library, \mathcal{T} , consists of 11 trajectories and is shown in Figure 3(a). The trajectories $x_i(t) : [0, T_i] \mapsto \mathbb{R}^4$ and the corresponding nominal open-loop control inputs were obtained via the direct collocation trajectory optimization method [3] by constraining $x_i(0)$ and $x_i(T_i)$ and locally minimizing a cost of the form:

$$J = \int_0^{T_i} [1 + u_0(t)^T R(t) u_0(t)] dt.$$

Here, R is a positive-definite cost matrix. For each $x_i(t)$ in \mathcal{T} we obtain controllers and funnels using the method described in Section IV. Similar to [31], we perform the verification on the time-varying nonlinear system by taking third-order Taylor-approximations of the dynamics about the nominal trajectories. For each trajectory, we use 11 sample points in time, t_i , for the verification. A 4.1 GHz PC with 16 GB RAM and 4 cores was used for the computations. The time taken for Step 1 of Algorithm 1 during one iteration of the alternation was approximately 10 seconds. Steps 2 and 3 take approximately 45-50 seconds each. Convergence is typically observed within 5 to 10 iterations of the algorithm. Three of the funnels in our library are shown in Figure 3(b). Note that the funnels have been projected down from the original

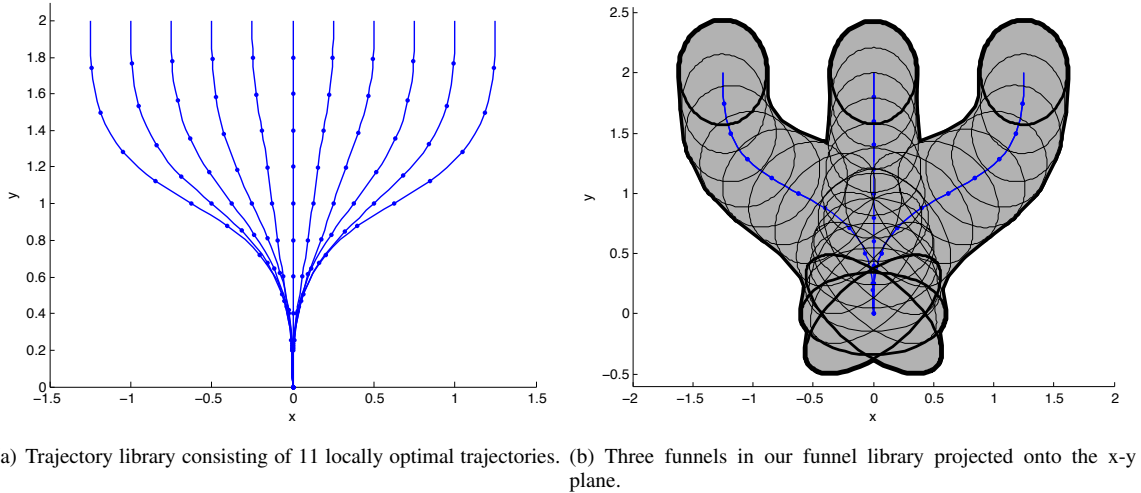


Fig. 3. Trajectory and funnel libraries for the plane.

four dimensional state space to the x-y plane for the sake of visualization.

Figure 4 demonstrates the use of the online planning algorithm in Section VI. The plane plans two funnels in advance while nominally attempting to fly in the y-direction and avoiding obstacles. The sensor range allows the plane to sense up to $5m$ ahead. The projection of the full sequence of funnels executed by the plane is shown in the figure. Figures 4(a) and 4(b) show the plane flying through the same forest with identical initial conditions. The only difference is that the cross-wind term is biased in different directions. In Figure 4(a), the cross-wind is primarily blowing towards the right, while in Figure 4(b), the cross-wind is biased towards the left. Of course, the planner is not aware of this difference, but ends up following different paths around the obstacles as it is buffeted around by the wind.

Finally, we demonstrate the utility of explicitly taking into account uncertainty in Figure 5. There are two obstacles in front of the plane. The two options available to the plane are to fly straight in between the obstacles or to bank right and attempt to go around them. If the motion planner didn't take uncertainty into account and simply chose to maximize the average distance to the obstacles, it would choose the trajectory that banks right and goes around the obstacles. However, taking the funnels into account leads to a different decision: going straight in between the obstacles is safer even though the distance to the obstacles is smaller. The utility of safety guarantees in the form of funnels is especially important when the margins for error are small and making the wrong decision can lead to disastrous consequences.

VIII. DISCUSSION AND FUTURE WORK

A. Stochastic Verification

Throughout this paper, we have assumed that all disturbances and uncertainty are bounded with probability one. In practice, this assumption may either not be fully valid or could lead to over-conservative performance. In such situations, it is more natural to provide guarantees of reachability of a




Fig. 5. This figure shows the utility of explicitly taking uncertainty into account while planning. The intuitively more “risky” strategy of flying in between two closely spaced obstacles is guaranteed to be safe, while the path that avoids going in between obstacles is less robust to uncertainty and could lead to a collision.

probabilistic nature. Recently, results from classical martingale theory have been combined with sums-of-squares programming in order to compute such probabilistic certificates of finite time invariance [29], i.e. provide upper bounds on the probability that a stochastic nonlinear system will leave a given region of state space. The results presented in [29] can be directly combined with the approach presented in this work to perform robust online planning on stochastic systems and will be the subject of future work.

B. Continuously Parameterized Families of Funnels

As discussed in Section V, we are currently partially exploiting invariances in the dynamics by shifting trajectories (and corresponding funnels) that we want to execute next in the cyclic coordinates so they line up with the cyclic coordinates of the robot's current state. In our example from Section VII, this simply corresponds to translating and rotating funnels so the

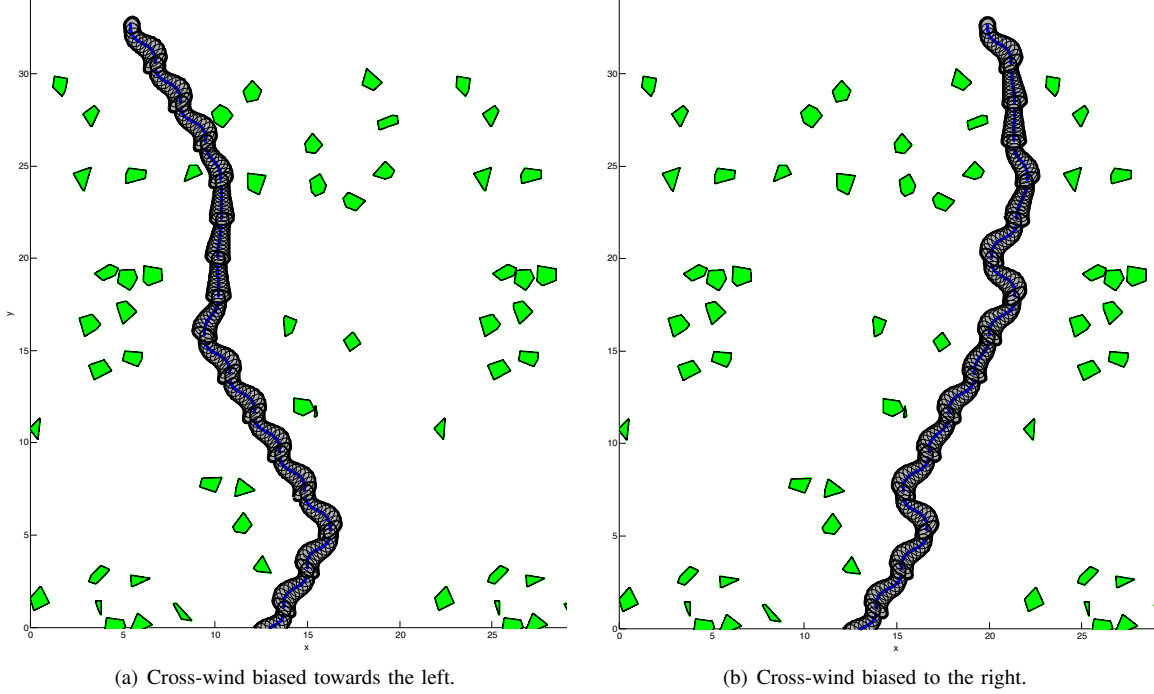


Fig. 4. Robust online planning through a forest of polygonal obstacles. The two subfigures show the plane flying through the same forest, but with the cross-wind biased in different directions (the planner is not aware of this difference). The eventual paths through the forest are significantly different, but the plane navigates the forest safely in each case.

beginning of the next trajectory lines up with the current state's x, y and yaw. However, we could further exploit invariances in the dynamics by shifting funnels around locally to ensure that they don't intersect an obstacle while still maintaining the current state inside the funnel. One can then think of the nominal trajectories and funnels being *continuously parameterized* by shifts in the cyclic coordinates. Interestingly, it is also possible to use sums-of-squares programming to compute conservative funnels for cases in which one shifts the nominal trajectory in the *non-cyclic* coordinates [20]. Thus, one could potentially significantly improve the richness of the funnel library by pre-computing continuously parameterized funnel libraries instead of just a finite family. However, choosing the right "shift" to apply at runtime is generally a non-convex problem (since the free-space of the robot's environment is non-convex) and thus one can only hope to find "shifts" that are locally optimal.

Another way to obtain continuously parameterized funnels is to ensure that *scalings* of $\rho(t)$ result in sub-level sets of $V(\bar{x}, t)$ that are invariant. This is equivalent to computing outer approximations of reachable sets for scalings of the initial condition set. Specifically, the SOS program (3) can be modified to guarantee that for some $\varepsilon > 0$, $\forall c \in [\varepsilon, 1]$ the sub-level sets defined by $c\rho(t)$ are invariant. One can visualize this as a continuously parameterized nested set of funnels. At runtime, this allows us to choose from scalings of our funnels. This could potentially reduce the number of funnels we need in our library. However, in practice, it may be difficult to obtain controllers and funnels that make guarantees of this form, especially if the disturbance terms are large.

C. Sequence optimization for Large Funnel Libraries

For extremely large funnel libraries, it may be computationally difficult to search all the funnels while planning online. This is a problem that traditional trajectory libraries also face [7]. Advances in submodular sequence optimization were leveraged in [7] to address this issue. The approach involves limiting the set of trajectories that are evaluated online and optimizing the sequence in which trajectories are evaluated. Guarantees are provided on the sub-optimality of the resulting strategy. This technique could be adapted to work in our framework too and will be addressed in future work.

D. Designing Funnel Libraries

One issue that we have not addressed in this paper is the choice of motion primitives in our library. While there has been considerable work on designing trajectory libraries (see Section II), designing *funnel libraries* poses challenges that go beyond just choosing a good set of nominal trajectories. The effect of uncertainty and feedback must be taken into account while constructing the library. One interesting problem domain in which it may be possible to design funnel libraries in a principled way is the case where the statistics of obstacle positions are known *a priori* (but the actual positions are unknown). An example of such a scenario is the task considered in this paper: autonomous UAV flight through a forest. It is known that the location of trees in a forest is well modeled by Poisson distributions [14]. Another example is legged robot locomotion on rough terrain, where the statistics of terrain variations are known beforehand. In such scenarios, it may be possible to design a randomized algorithm in the spirit of LQR-Trees [31] where one attempts to plan paths through particular

realizations of the environment by sequencing funnels together and adding a funnel to the library every time a collision free sequence of funnels is not found in the existing library. Under certain assumptions on the distributions of obstacles (e.g. stationarity, ergodicity), it is conceivable that such a randomized algorithm may be probabilistically complete.

IX. CONCLUSION

In this paper, we have presented an approach for motion planning in *a priori* unknown environments with dynamic uncertainty in the form of bounded parametric model uncertainty, disturbances, and state errors. The method augments the traditional trajectory library approach by constructing stabilizing controllers around the nominal trajectories that explicitly attempt to minimize the size of the reachable set of the system subjected to disturbances and uncertainties. The precomputed set of reachable sets (“funnels”) is then used to plan online by *sequentially composing* them together in a manner that ensures obstacles are avoided. By explicitly taking into account uncertainty and disturbances while making motion plans, we can evaluate trajectory sequences based on how susceptible they are to disturbances. We have demonstrated our approach on a simulation of a plane flying in two dimensions through a forest of polygonal obstacles. Future work will focus on generating funnel libraries automatically for environments with known obstacle distributions (e.g. forests) and extending our results to scenarios in which a stochastic description of uncertainty is more appropriate.

ACKNOWLEDGMENT

This work was partially supported by ONR MURI grant N00014-09-1-1051 and the MIT Intelligence Initiative. Anirudha Majumdar is currently partially supported by a Siebel Foundation Scholarship.

REFERENCES

- [1] A. A. Ahmadi. Non-monotonic Lyapunov Functions for Stability of Nonlinear and Switched Systems: Theory and Computation. Master's thesis, Massachusetts Institute of Technology, June 2008.
- [2] D. Berenson, R. Diankov, K. Nishiwaki, S. Kagami, and J. Kuffner. Grasp planning in complex scenes. In *International Conference on Humanoid Robots*, pages 42–48. IEEE, 2007.
- [3] J. T. Betts. *Practical Methods for Optimal Control Using Nonlinear Programming*. SIAM Advances in Design and Control. Society for Industrial and Applied Mathematics, 2001.
- [4] A. Bry and N. Roy. Rapidly-exploring random belief trees for motion planning under uncertainty. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Shanghai, China, 2011.
- [5] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek. Sequential composition of dynamically dexterous robot behaviors. *International Journal of Robotics Research*, 18(6):534–555, June 1999.
- [6] E. F. Camacho and C. Bordons. *Model Predictive Control*. Springer-Verlag, 2nd edition, 2004.
- [7] D. Dey, T. Liu, B. Sofman, and D. Bagnell. Efficient optimization of control libraries. Technical report, Technical Report (CMU-RI-TR-11-20), 2011.
- [8] E. Frazzoli, M. Dahleh, and E. Feron. Maneuver-based motion planning for nonlinear systems with symmetries. *IEEE Transactions on Robotics*, 21(6):1077–1091, December 2005.
- [9] E. Frazzoli, M. A. Dahleh, and E. Feron. Real-Time Motion Planning for Agile Autonomous Vehicles. *Journal of Guidance, Control, and Dynamics*, 25(1):116–129, January/February 2002.
- [10] J. Gillula, H. Huang, M. Vitus, and C. Tomlin. Design of guaranteed safe maneuvers using reachable sets: Autonomous quadrotor aerobatics in theory and practice. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1649–1654. IEEE, 2010.
- [11] C. Green and A. Kelly. Toward optimal sampling in the space of paths. In *13th International Symposium of Robotics Research*, November 2007.
- [12] P. Jacobs and J. Canny. Robust motion planning for mobile robots. In *Robotics and Automation, 1990. Proceedings., 1990 IEEE International Conference on*, pages 2–7. IEEE, 1990.
- [13] A. A. Julius and G. J. Pappas. Trajectory based verification using local finite-time invariance. In *HSCC '09: Proceedings of the 12th International Conference on Hybrid Systems: Computation and Control*, pages 223–236, Berlin, Heidelberg, 2009. Springer-Verlag.
- [14] Karaman and Frazzoli. High-speed flight through an ergodic forest. In *IEEE Conference on Robotics and Automation*, 2012.
- [15] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *Int. Journal of Robotics Research*, 30:846–894, June 2011.
- [16] J. Kuffner and S. Lavalle. RRT-connect: An efficient approach to single-query path planning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 995–1001, 2000.
- [17] C. Liu and C. Atkeson. Standing balance control using a trajectory library. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 3031–3036. IEEE, 2009.
- [18] C. Maier and F. Allgower. A set-valued filter for discrete time polynomial systems using sum of squares programming. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pages 223–228. IEEE, 2009.
- [19] A. Majumdar, A. A. Ahmadi, and R. Tedrake. Control design along trajectories with sums of squares programming. *Under review*, 2013.
- [20] A. Majumdar, M. Tobenkin, and R. Tedrake. Algebraic verification for parameterized motion planning libraries. In *Proceedings of the 2012 American Control Conference (ACC)*, 2012.
- [21] D. Mayne, M. Seron, and S. Rakovic. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224, 2005.
- [22] J. L. Ny and G. Pappas. Sequential composition of robust controller specifications. In *International Conference on Robotics and Automation (ICRA)*. IEEE, 2012.
- [23] P. A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, May 18 2000.
- [24] R. Platt, R. Tedrake, L. Kaelbling, and T. Lozano-Perez. Belief space planning assuming maximum likelihood observations. In *Proceedings of Robotics: Science and Systems*, 2010.
- [25] P. Reist and R. Tedrake. Simulation-based LQR-trees with input and state constraints. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2010.
- [26] T. Schouwenaars, B. Mettler, E. Feron, and J. How. Robust motion planning using a maneuver automation with built-in uncertainties. In *American Control Conference, 2003. Proceedings of the 2003*, volume 3, pages 2211–2216. IEEE, 2003.
- [27] P. Sermanet, M. Scoffier, C. Crudele, U. Muller, and Y. LeCun. Learning maneuver dictionaries for ground robot planning. In *Proc. 39th International Symposium on Robotics (ISR08)*, 2008.
- [28] A. Shkolnik. *Sample-Based Motion Planning in High-Dimensional and Differentially-Constrained Systems*. PhD thesis, MIT, February 2010.
- [29] J. Steinhardt and R. Tedrake. Finite-time regional verification of stochastic nonlinear systems. In *Proceedings of Robotics: Science and Systems (RSS) 2011*, January 17 2011.
- [30] M. Stolle and C. Atkeson. Policies based on trajectory libraries. In *International Conference on Robotics and Automation (ICRA)*. IEEE, 2006.
- [31] R. Tedrake, I. R. Manchester, M. M. Tobenkin, and J. W. Roberts. LQR-Trees: Feedback motion planning via sums of squares verification. *International Journal of Robotics Research*, 29:1038–1052, July 2010.
- [32] M. M. Tobenkin, I. R. Manchester, and R. Tedrake. Invariant funnels around trajectories using sum-of-squares programming. *Proceedings of the 18th IFAC World Congress, extended version available online: arXiv:1010.3013 [math.DS]*, 2011.