Randomized Model Predictive Navigation

Herbert G. Tanner[†], {btanner@udel.edu}

Jorge L. Piovesan[‡]{jpiovesan@ka-wireless.com}

[†]*Dept. of Mechanical Engineering, University of Delaware*
Newark DE, 19716
[‡]*K&A Wireless LLC, 2617 Juan Tabo NE Suite A, Albuquerque NM 87112*

UDME Technical Report
Number UDMETR-2010-0001
February 2010

Department of Mechanical Engineering
University of Delaware
Newark, Delaware
URL: http://www.me.udel.edu

Randomized Model Predictive Navigation

Herbert G. Tanner[†], {btanner@udel.edu}

Jorge L. Piovesan[‡]{jpiovesan@ka-wireless.com}

[†]Dept. of Mechanical Engineering, University of Delaware
Newark DE, 19716
[‡]K&A Wireless LLC, 2617 Juan Tabo NE Suite A, Albuquerque NM 87112

University of Delaware, Dept. of ME, UDMETR-2010-0001

February 2010

Abstract

A new methodology for implementing nonlinear receding horizon optimization is presented, with direct application to robot navigation in cluttered environments. The methodology combines elements from statistical learning theory with nonlinear receding horizon schemes that use control Lyapunov functions as terminal costs, while relaxing the conditions on the time derivatives of the latter, based on a new result for stability of nonlinear systems with switching dynamics. As the theoretical analysis indicates, and numerical results verify, the proposed receding horizon scheme can utilize terminal costs that are not control Lyapunov functions. The resulting strategy is shown to outperform traditional potential field-based techniques, even when additional optimization objectives are imposed, and allows for trade-offs between performance and computational complexity.

1 Introduction

1.1 Overview

The discussion in this paper focuses on robotic systems with switching dynamics which can be expressed in the form

$$\dot{x} = f_\sigma(x,u) \quad , \tag{1}$$

with $x \in \mathbb{R}^n$ being the state, $u \in \mathbb{R}^m$ the control input, and $\sigma \in \Sigma \subseteq \mathbb{R}$ being the index that indicates the particular dynamics currently active. It is assumed that the system's workspace is a set $\mathscr{W} \subset \mathbb{R}^n$ and that the system is constrained to stay clear of a possibly disconnected region of the state space $\mathscr{O} \subset \mathscr{W}$ which may represent areas occupied by obstacles or sets of states that violate certain constraints. Let $Q \triangleq \mathscr{W} \setminus \mathscr{O}$ denote the remaining, admissible region of the system's workspace. The boundary of Q is assumed known exactly.

The problem is to design a feedback control law $u = k(x)$ to steer to, and asymptotically stabilize the system at, a desired state x_d which for simplicity it is assumed to be the origin in $\mathbb{R}^n$, along trajectories that belong in Q for all time, and are optimal (or near optimal in the sense of [2]) in terms of some task objectives. The control law does not have to be continuous in x, nor time-invariant, but it will be assumed to be piecewise continuous in x and continuous in time.

Problems of this nature arise when mobile robots with limited power resources need to navigate in cluttered environments, without reliance on pre-computed motion plans. A justification for the latter requirement could be the presence of uncertainty, either in the system dynamics or in the description of the environment, and possibly the need for re-tasking and reconfiguration, both of which necessitate the use of real-time feedback control signals. These problems are slightly more challenging than standard robot navigation problems, particularly due to the desire for on-line computed optimal control strategies.

The approach outlined in this paper is a blend of control Lyapunov function-based model predictive control schemes, with navigation functions and randomized optimization algorithms. What motivated the search for a solution with these ingredients is the desire for real-time computation by miniature mobile robots of optimal or near-optimal motion plans with guaranteed stability and convergence properties. Ideally, although not addressed in this paper, the control design scheme should be robust to environment and model uncertainty. A nonlinear receding horizon architecture is chosen because it allows the real-time generation of sub-optimal control laws for systems of the form (1). It has been shown that such architectures, when combined with control Lyapunov functions (CLFs), can guarantee closed loop asymptotic stability [3]. The motion planning nature of the task, as well as the need for a control Lyapunov function, suggested the integration of navigation functions into the nonlinear receding horizon framework. Finally, the role of randomized algorithms is to simplify the online optimization task, enforce some task-specific constraints, and offer the designer a clear path to robustifying the overall architecture against model and environment uncertainty. The latter task is part of ongoing work, and research directions along these lines are outlined in Section 6.

1.2 Contribution and related work

It has been demonstrated [4, 3, 5] that the use of control Lyapunov functions as terminal costs in receding horizon nonlinear optimization offers closed loop stability in the infinite horizon, as well as robustness [6]. These features are reported as improvements over earlier approaches which enforced stability by means of terminal equality constraints [7], more so since the latter formulation has an adverse effect on computational complexity. For the unconstrained case, it is shown [8, 6] that an appropriately long control horizon allows the terminal cost to be set to zero, although numerical conditioning considerations may dictate that a nonzero cost is used nonetheless [3]. Practically, the solution of nonlinear receding horizon problems requires the use of numerical optimization packages, among which non-proprietary examples include RIOTS [9] and NTG [10].

Randomized algorithms have been suggested as a viable alternative to exact analytic or numerical optimization techniques [11, 12]. Using randomized algorithms for control synthesis was motivated by early results on the complexity of several robust stability design problems [13], and following the categorization of [14], existing approaches fall into two groups; the first is based primarily on the Vapnik-Chervonenkis theory of statistical learning, and representative samples include [2, 15]; the second employs sequential approximation techniques as in [14, 16, 17, 18, 19]. Part of the elegance of these randomized approaches is due to the fact that they offer explicit bounds on the number of samples required to guarantee specific performance levels. This set of statistical learning tools have proven effective in linear robust control design problems [12, 2].

Robot motion planning in the presence of obstacles is also an NP-hard problem [20]. To address such problems both purely analytical [21], numerical [22], randomized [23], as well as hybrid [24] methods have been proposed, with the latter aimed at addressing through graph search techniques the problem of local minima present in the classical potential approach of [21]. Despite this limitation, potential fields have been particularly appealing due to their ability to provide feedback motion planning controllers, and several different solutions have been proposed to the local minima problem [25, 26, 27, 28]. Of the latter, probably the most popular technique being the approach of navigation functions [28], a special potential function construction which can be tuned[1] so that the only minimum is the desired configuration. From a control theoretic perspective, an additional desirable feature of a navigation function is that it serves as a natural Lyapunov function candidate. For time-varying or switching systems, standard Lyapunov stability conditions may impose limitations on the use of navigation functions, since the gradient of latter cannot be guaranteed to be non-vanishing, except for trivial cases.

The technical challenge addressed in this paper is that of integrating nonlinear model predictive control, navigation function-based motion planning, and randomized algorithms, while preserving the desirable attributes of the original methods, such as the feedback character of the control laws, the closed loop stability, the sub-optimality of the solutions, and the guaranteed constraint satisfaction. To achieve this integration, the paper suggests a new set of sufficient conditions for asymptotic stability in switching nonlinear systems, which are relaxed compared to related results that involve a common Lyapunov function, in the sense that the function does not need to have a negative semidefinite time derivative. Based on this result, the stability properties of CLF-based receding horizon schemes are re-established along the lines of the analysis in [3]. The last part of the paper's technical discussion involves a class of randomized algorithms used in distribution-free statistical learning methods, in order to solve the finite horizon nonlinear optimal control problem. The method employed here falls within the class of the VC-theory based methods, with the choice motivated by the fact that at this stage of this work uncertainty is not directly addressed. The integration of this component relies on parameterization of the set of solutions to the model predictive control design problem, and subsequent

[1] Although theoretical results guarantee the existence of lower bounds for the admissible values of the tuning parameters, the tuning process might be challenging, especially in high dimensional spaces.

sampling of the parameter space using a number of samples sufficiently large to guarantee an approximation of the optimal solution to a desired level of confidence.

From the theoretical standpoint, the paper contributes by relaxing the requirements on the terminal cost as they appear in [3], and allow the class of terminal cost candidates to include a wider range of positive definite functions —not necessarily control Lyapunov functions. From the applications standpoint, the proposed approach introduces a new set of (randomized) optimization tools within the receding optimization framework, and links this type of control design methodology to problems in robot navigation in the presence of constraints. In the proposed mix, the technologies involved complement each other: Navigation functions offer constrained finite horizon optimal solutions without the explicit use of constraint conditions in optimization. The receding horizon framework alleviates the problem of navigation function parameter tuning by enhancing the field of view compared to the "myopic" gradient descent (some local minima can be tolerated). The randomized algorithms facilitate the implementation of receding horizon schemes by eliminating the dependence on specialized numerical optimization libraries, and offer trade-offs between control performance and computation speed. In addition, due to the randomized approach to the receding horizon optimization, the resulting architecture does not rely on a particular model structure: the system dynamics just needs to be considered in the formulation of the terminal cost function, which when constructed as a navigation function offers a vehicle to incorporate state constraints without explicitly stating them as such during optimization. Quite recently, a "dual," hierarchical approach was explored [29], where state constraints are enforced at the higher level using a model predictive scheme while collision free trajectory generation is done using navigation functions. Here, navigation functions are used as a vehicle for efficiently introducing state constraints into the receding horizon framework, and simultaneously establishing stability for the infinite horizon.

1.3 Organization

Section 2 that follows discusses stability issues related to using navigation functions as Lyapunov functions in systems with switching dynamics, and suggests a remedy for cases where switching is not arbitrarily fast. Section 3 integrates navigation functions into a receding horizon control scheme in the role of terminal cost functions, and establishes the asymptotic stability of the closed loop system by exploiting the results of the preceding section. The nonlinear optimization problem is then approached from a statistical learning tool perspective in Section 4, where the computation of probable near optimal control laws is discussed.

2 Navigation Functions and Switching Dynamics

Navigation functions have been long used in robot motion planning, and although not always explicitly used as such, they are natural control Lyapunov function candidates. This section addresses some issues related to the application of navigation functions for motion planning and control in systems with switching dynamics. The discussion highlights the fact that switching system stability requires more stringent conditions which navigation functions may not readily satisfy. It is thus deemed useful to develop alternative stability requirements, which relax some of the assumptions placed by existing techniques. As expected, there is a trade-off, and in this case this comes in the form of restrictions on the rate of change of the control Lyapunov function as well as on switching times.

A navigation function on $\varphi : Q \to [0,1]$, where Q denotes the obstacle-free configuration space of a robot, is a function that satisfies the following properties [28, 30]:

1. is smooth (or at least C^2),
2. has a unique minimum at a single point $q_d \in Q$,
3. is admissible (uniformly maximal) on the boundary of Q, and
4. is a Morse function (its critical points are non-degenerate).

It has been shown [31] that the function

$$\varphi(q) = \frac{\|q - q_d\|^2}{\left(\|q - q_d\|^{2\kappa} + \beta\right)^{1/\kappa}} \quad ,$$

where β is the product of a number of functions $\beta_i = \|q - q_i\| - r_i$, each representing a spherical obstacle of radius r_i centered at q_i in the robot's workspace, can be made into a navigation function for a sufficiently large value of $\kappa > 0$. The product includes a boundary obstacle function expressed as $\beta_0 = r_0^2 - \|q - q_0\|^2$, with r_0 being the radius of the robot's spherical workspace, centered at configuration q_0. The effect of increasing κ is that obstacles appear "sharper" in the potential function, and their influence is diminished away from their boundary. The problem with increasing κ is that the free space becomes "flatter," gradients are small and scaling them into velocity reference commands could become problematic, and the effect of obstacles may be felt only when the robot is too close to them.

2.1 Control via navigation functions

The idea behind potential field methods is to steer the system at hand along the direction of the negated gradient of the navigation function. For cases where the kinematics can be trivially described by $\dot{q} = u$, this is achieved by setting $u = -\nabla_q \varphi(q)$. In that sense, one can consider the navigation function being used as a control Lyapunov function. A control Lyapunov function [32] is a differentiable, positive definite and proper (radially increasing) function of the state $V : \mathrm{R}^n \to \mathrm{R}_+$ for which

$$\inf_u \left\{ \frac{\partial V(x)}{\partial x} f(x,u) \right\} < 0, \quad \forall x \neq 0 \ . \tag{2}$$

For the case of navigation function-based feedback of the form $u(q) = -\nabla_q \varphi(q)$, if $\varphi(q)$ is used as a control Lyapunov function, one obtains $\frac{\partial V(q)}{\partial q} f(q) = -\|\nabla_q \varphi(q)\|$, and (2) is satisfied almost everywhere (*a. e.*) in Q, with the exception of a set of measure zero containing the saddle points of the navigation function. These points might appear as a small nuisance in the case of continuous-state feedback control, but they present a more significant challenge in the case of discontinuous feedback, since then the system is essentially time-varying and (2) has to be strengthened in the form [33, 34, 35]

$$\inf_u \left\{ \frac{\partial V(x)}{\partial x} f(x) \right\} \leq -W(x) \ , \tag{3}$$

where $W(x)$ is a continuous positive definite function of the state. Condition (3) in particular, cannot be satisfied when navigation functions are used as control Lyapunov functions, unless one is prepared to exclude neighborhoods of the saddle points. Similar stability considerations may arise in the case of time-invariant feedback stabilization of nonholonomic systems, particularly when robustness to measurement errors is desired [36].

This motivates one to seek alternative conditions for asymptotic stability in the case of switching dynamics or discontinuous feedback, which are relaxed compared to (3) in the sense that they do not bound the derivative of the control Lyapunov function below another negative definite function.

2.2 Lyapunov-based extensions

This section suggest a set of conditions for asymptotic stability of switching systems for which switching intervals lengths can be lower bounded by arbitrarily small bounds. The conditions presented allow a Lyapunov-like function to temporarily increase in a controlled fashion along the trajectories of the switched system, if it can be ensured that by the next switching instant the function will have been sufficiently decreased. Thus, the time derivative of this function is no longer required to be negative definite, but in the case it is, one may allow the switching interval lengths to collapse to zero and recover the known conditions for uniform asymptotic stability.

The conditions placed by the following proposition can be hard to verify in the general case, and therefore cannot substitute for the classical sign definiteness inequalities. One, however, can find instances of control design problems (as in the case of receding horizon control) when these conditions may be enforced, either constructively or in the form of (optimization) constraints.

Definition 1 (Nonzero Interval Switching sequence)**.** *A sequence* $\mathbf{s} \triangleq \{(\tau_k, \sigma_k)\}_{k=0}^{\infty}$*, where* σ_k *is an index from a set* $\Sigma \subseteq \mathbb{R}$*, is a nonzero interval switching sequence if there exists an arbitrarily small* $\underline{\tau} > 0$ *such that* $\tau_{n+1} - \tau_n \geq \underline{\tau}$ *for all* $n \in \mathbb{N}$*.*

Proposition 1. *Let $V(x)$ be a differentiable, positive definite function that is lower bounded by a class-$\mathscr{K}_\infty$ function $\theta(\|x\|)$. If for a class-$\mathscr{K}_\infty$ function $\gamma(\cdot)$, within each of the switching intervals $(\tau_i, \tau_{i+1}]$ in* $\mathbf{s}$

$$\max_{t\in(\tau_i,\tau_{i+1}]} \left\{ \frac{\partial V}{\partial x} f(x) \right\} < \frac{\theta(\|x(\tau_i)\|)}{\tau_{i+1}-\tau_i} \ , \tag{4}$$

$$\int_{\tau_i}^{\tau_{i+i}} \frac{\partial V}{\partial x} f(x)\, dt \le -\gamma(\|x(\tau_i)\|) \ , \tag{5}$$

then (1) *is asymptotically stable at the origin.*

Proof. Pick an $\varepsilon > 0$ and define a ball $B_\varepsilon \triangleq \{x \in \mathbb{R}^n \mid \|x\| \le \varepsilon\}$ (Figure 1). On the boundary of B_ε, $V(x)$ attains a minimum by continuity. Set $V_\varepsilon \triangleq \min_{x\in\partial B_\varepsilon} V(x)$ and define Ω_ε to be the path connected component of the set $\{x \in \mathbb{R}^n \mid V(x) \le V_\varepsilon\}$, which contains the origin. Let $V_\delta \triangleq V_\varepsilon - \theta(\|x(\tau_i)\|) > 0$ and similarly define Ω_δ as the path connected component of the set $\{x \in \mathbb{R}^n \mid V(x) \le V_\delta\}$ containing the origin. Fit a ball B_δ inside Ω_δ, by defining it as $B_\delta = \{x \in \mathbb{R}^n \mid \|x\| \le \min_{x\in\partial\Omega_\delta} \|x\|\}$. Now let $\Omega_d \triangleq \{x \in \mathbb{R}^n \mid V(x) \le V_\delta - \gamma(\|x(\tau_i)\|)\}$. Assume that at time τ_i,

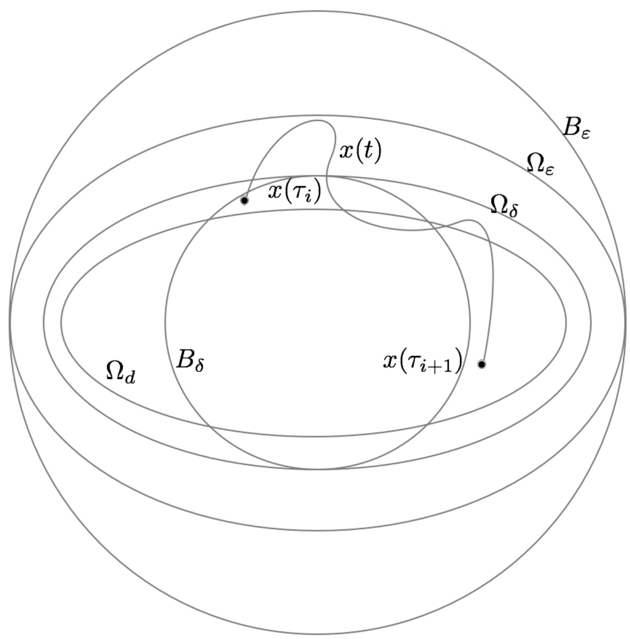

Figure 1: The sets defined in the proof of Proposition 1.

$x(\tau_i) \in B_\delta$. Since $B_\delta \subseteq \Omega_\delta$, $V(x(\tau_i)) \le V_\delta = V_\varepsilon - (\tau_{i+1} - \tau_i)d$. For the evolution of $V(x(t))$ in $t \in (\tau_i, \tau_{i+1})$ one has $V(x(t)) = V(x(\tau_i)) + \int_{\tau_i}^{t} \frac{\partial V}{\partial x} f_\sigma(x(\tau))\, d\tau$ and due to (4), $V(x(t)) \le V(x(\tau_i)) + d(t - \tau_i)$, where $d \triangleq \frac{\theta(\|x(\tau_i)\|)}{\tau_{i+1}-\tau_i}$. Thus, $x(t) \in \Omega_\varepsilon$, for all $t \in (\tau_i, \tau_{i+1}]$. In addition, due to (5) $x(\tau_{i+1}) \in \Omega_\delta$, and an inductive argument using (4) and (5) shows that $x(\tau_k) \in \Omega_\delta \subset \Omega_\varepsilon$ for all $k \ge i$. Thus stability is established. Note though that the level sets of $V(x)$ are not necessarily invariant: even if $x(\tau_i) \in \Omega_\delta$, $x(t)$ for $t \in (\tau_i, \tau_{i+1}]$ is not forced to stay in Ω_δ as in the case where $\dot{V}(x)$ is negative semidefinite.

Condition (5) implies the existence of a converging sequence $\{V\big(x(\tau_k)\big)\}$, since $V(x)$ is lower bounded by 0 away from $x = 0$ and it is strictly decreasing. It follows that as $k \to \infty$, the difference between two consecutive terms $\|V\big(x(\tau_i)\big) - V\big(x(\tau_{i+1})\big)\|$ should converge to zero. Unless $\lim_{k\to\infty} V\big(x(\tau_k)\big) = 0$, this means that $\gamma(\|x(\tau_k)\|) \to 0$, without $\|x(\tau_k)\| \to 0$, which is impossible since $\gamma(\cdot)$ is a class-$\mathscr{K}_\infty$ function. Therefore, $\{V\big(x(\tau_k)\big)\} \to 0$ as $k \to \infty$ which suggests that $\|x(\tau_k)\| \to 0$ when $k \to \infty$. As a result, (1) is asymptotically stable at the origin. □

Contrary to existing asymptotic stability conditions, the ones of Proposition 1, although more relaxed, do not appear particularly useful for they may not be verifiable in a straightforward manner: the switching sequence needs to be known a priori, and then the function's directional derivative has to be integrated along the system's trajectories to check (5). Nonetheless, there are instances of receding horizon control design where these relaxations are applicable. This is particularly true when the receding horizon optimization is performed in a randomized way, due to the fact that (5) may be enforced by construction. The next section binds Proposition 1 to the CLF-based receding horizon framework and verifies that asymptotic stability for the infinite horizon is recovered even under the relaxations of Proposition 1.

3 Model Predictive Navigation

In this section that a function that satisfies the requirements of Proposition 1 can be used as a terminal cost in a receding horizon scheme, to ensure stability in the infinite horizon without the need to impose additional terminal cost constraints on the optimization process.

In finite horizon optimization, one aims to minimize the functional

$$J_T\big(x,u(\cdot)\big) \triangleq \int_0^T q\big(x^u(\tau;x),u(\tau)\big)\,\mathrm{d}\tau + V\big(x^u(T;x)\big) \quad , \tag{6}$$

that quantifies the cost of flowing along a closed loop system trajectory

$$x^u(t;x) \triangleq x + \int_0^t f\big(x^u(\tau;x),u(\tau)\big)\,\mathrm{d}\tau \quad , \tag{7}$$

starting at x, under the control law $u(\tau)$, for a time interval in which $t \in [0,T]$. The function $V(\cdot)$ in (6) is an approximation to the tail of the truncated integral compared to the infinite horizon case where

$$J_\infty\big(x,u(\cdot)\big) \triangleq \int_0^\infty q\big(x^u(\tau;x),u(\tau)\big)\,\mathrm{d}\tau \quad .$$

The integrand $q : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}_+$ is a C^2 function that measures the performance of the system, by penalizing both state and control, and just as in [3] it is assumed to satisfy

$$q(x,u) \geq c_q(\|x\|^2 + \|u\|^2), \quad x \in \mathbb{R}^n,\ u \in \mathbb{R}^m$$

for some $c_q > 0$ and $q(0,0) = 0$, and be a convex function $u \mapsto q(x,u)$ for each $x \in \mathbb{R}^n$.

The problem is to determine a control law that minimizes the finite horizon cost from x,

$$u^*_T(x;\cdot) \triangleq \arg\min J_T\big(x,u(\cdot)\big) \quad ,$$

giving rise to an optimal finite horizon trajectory $\big(x^*_T(t;x),u^*(t;x)\big)$, for $t \in [0,T]$, when $u^*(x;\cdot)$ is used in (7) to compute $x^*_T(t;x) = x + \int_0^T f\big(x^*_T(\tau;x),u^*_T(\tau)\big)\,\mathrm{d}\tau$. In a receding horizon strategy, one uses $u^*_T(t;x)$ for $t \in [0,\zeta]$, with $\zeta < T$, and then recomputes $u^*\big(t;x^*(\zeta,x)\big)$. One issue related to this strategy is whether the sequence of finite horizon optimal control laws $u^*_T(t;x_k)$, with $x_k = x\big(\sum_{i=1}^k \zeta_i\big)$ stabilizes (1) for the infinite horizon, that is for $t \in [0,\infty)$. For it can be demonstrated that although $J^*_T(x) \triangleq J_T\big(x,u^*(\cdot)\big)$ can behave nicely, the closed loop system may be unstable, particularly so in a case that combines setting V in (6) to zero with a small prediction horizon T.

Control Lyapunov functions offer a solution to the infinite horizon stability problem of receding horizon control schemes [4, 5] when used to approximate the terminal cost within a finite horizon optimization. Although relaxed stability conditions have been suggested [6], the standard assumption for the control Lyapunov function in this framework [3] is that it is compatible with the incremental cost $q(x,u)$, in the following sense

$$\inf_{u(\cdot)} \left(\frac{\partial V}{\partial x} f(x,u) + q(x,u) \right) \leq 0 \quad . \tag{8}$$

This section suggests that the use of a navigation function as a control Lyapunov function in the role of the terminal cost V is sufficient to provide stability for the closed loop system. Such a choice appears appealing because the navigation function becomes a vehicle for introducing state constraints into an (otherwise unconstrained) optimization problem, without stating them explicitly as such —doing so increases the computation burden significantly [3]. However, a navigation function cannot satisfy (8) in any small neighborhood of its critical (saddle) points. This is where

Proposition 1 is utilized. In this case the switching sequence is $\{\zeta_k\}_{k=1}^{\infty}$, and there are practical reasons (computation time for optimization) that force ζ_k to be nonzero, at any computational platform, no matter how fast.

The rest of the section traces the steps of the stability analysis in [3], showing why a terminal cost $V(x) \equiv \varphi(x)$ that is not formally a control Lyapunov function, and it is not monotonically decreasing as (8) is requiring, can be used in (6) with guaranteed infinite horizon stability, if the conditions of Proposition 1 are met. The following assumption is made in [3], and is adapted here by defining Ω as the path connected component of $\{x \in \mathbb{R}^n \mid 0 \leq V(x) \leq c_V < \infty\}$ containing the origin, where c_V is a positive constant. (If $V(x)$ is a navigation function, $c_V = 1$.)

Assumption 1 (Bounds on optimal cost)**.** *For all $x \in \Omega$ and $T > 0$ there exist positive constants c_T and m_∞*

1. $J^*_\infty(x) \geq m_\infty \|x\|^2$,
2. $J^*_T(x) \leq c_T \|x\|^2$.

The next assumption is the equivalent of (8) in [3], adapted to the relaxed requirements of Proposition 1.

Assumption 2 (Cost compatibility)**.** *The incremental cost and the terminal cost function are compatible in the sense that there exists a feedback control $u = k(x)$, a $\mathscr{K}_\infty$ class function $\gamma(\cdot)$, and a constant $\zeta > 0$ for which*

$$\int_0^{\zeta} q\big(x(\tau), k(x(\tau))\big)\, \mathrm{d}\tau \leq \gamma\big(\|x(0)\|\big) \leq -\int_0^{\zeta} \dot{V}\big(x(\tau), k(x(\tau))\big)\, \mathrm{d}\tau,\ \forall x(0) \in \Omega \ . \tag{9}$$

The left hand side of (9) can always be satisfied for a suitable choice of ζ: one can always find a class $\mathscr{K}$ function to bound locally from above the integral of the closed loop incremental cost, since the latter is a positive definite function of x. If $V\big(x, k(x)\big)$ is a Lyapunov function, then the right hand side of the inequality is trivially satisfied. Although a positive scaling on V might help separating the left and right parts of (9), the remaining challenge is the behavior of the corresponding bounds in the neighborhood of $x(0) = 0$. For example, if the left hand side of (9) is bounded by a linear function of $\|x(0)\|$, while the right hand side is bounded by a quadratic one, then these two bounds cannot reconcile. At the limit, when $\zeta \to 0$, (9) reduces to $(\dot{V} + q)(x, k(x)) \leq 0$ of [3], in which case one may set $\gamma(\|x\|) \triangleq q(x, k(x))$. When $\dot{V}$ is not negative semidefinite, one gets yet another shot at establishing stability if able to identify appropriate ζ and $k(x)$ such that both inequalities are satisfied.

Next comes a restatement of [3, Theorem 2.2.1], for a terminal cost function that satisfies (4) and (9). Note that although (4) is not explicitly invoked in the following discussion, it has to be enforced to ensure stability as well as the satisfaction of any constraints encoded in $V(x) \equiv \varphi(x)$. Specifically, $\theta(\|x_0\|)$ must be small enough so that $\varphi\big(x(t)\big) + \theta(\|x_0\|) < 1$, for all $t \in [0, \zeta]$.

Proposition 2. *Suppose that $x \in \mathbb{R}^n$ and $T > 0$ are such that*

$$x^*_T(T; x) \in \Omega \ . \tag{10}$$

Then, for any $\zeta(x)$, and feedback control $u = k(x(t))$ for which there exists a $\mathscr{K}_\infty$ class function γ satisfying (9)*:*

$$\int_T^{T+\zeta} \dot{V}(x(\tau), k(x(\tau)))\, \mathrm{d}\tau \leq -\gamma(\|x(T)\|) \leq -\int_T^{T+\zeta} q\big(x(T+\tau), k(x(T+\tau))\big)\, \mathrm{d}\tau \ ,$$

for $x \in \Omega$, the optimal cost from $x^*_T(\zeta; x)$ is such that

$$J^*_T(x^*_T(\zeta; x)) \leq J^*_T(x) - \int_0^{\zeta} q(x^*_T(\tau; x), u^*_T(\tau; x))\, \mathrm{d}\tau \ . \tag{11}$$

Proof. Let $(\tilde{x}(t), \tilde{u}(t))$, $t \in [0, T+\zeta]$, be the trajectory obtained by concatenating $(x^*_T, u^*_T)(t; x)$, $t \in [0, T]$, and $(x^k, u^k)(t-$

$T;x^*_T(T;x))$, $t\in[T,T+\zeta]$. Consider the cost of using $\tilde{u}(\cdot)$ for T seconds beginning at an initial state $x^*_T(\zeta;x)$,

$$
\begin{aligned}
J_T(x^*_T(\zeta;x),\tilde{u}(\cdot)) &= \int_{\zeta}^{T+\zeta} q(\tilde{x}(\tau),\tilde{u}(\tau))\,\mathrm{d}\tau + V(\tilde{x}(T+\zeta))\\
&= \int_{\zeta}^{T} q(x^*(\tau;x),u^*(\tau;x))\,\mathrm{d}\tau + \int_{T}^{T+\zeta} q(\tilde{x}(\tau),\tilde{u}(\tau))\,\mathrm{d}\tau + V(\tilde{x}(T+\zeta))\\
&= J^*_T(x) - \int_0^T q(x^*_T(\tau;x),u^*_T(\tau;x))\,\mathrm{d}\tau - V(x^*_T(T;x)) + \int_{\zeta}^{T} q(x^*(\tau;x),u^*(\tau;x))\,\mathrm{d}\tau\\
&\quad + \int_{T}^{T+\zeta} q(\tilde{x}(\tau),\tilde{u}(\tau))\,\mathrm{d}\tau + V(\tilde{x}(T+\zeta))\\
&= J^*_T(x) - \int_0^{\zeta} q(x^*_T(\tau;x),u^*_T(\tau;x))\,\mathrm{d}\tau - V(x^*_T(T;x)) + \int_{T}^{T+\zeta} q(\tilde{x}(\tau),\tilde{u}(\tau))\,\mathrm{d}\tau + V(\tilde{x}(T+\zeta))\\
&= J^*_T(x) - \int_0^{\zeta} q(x^*_T(\tau;x),u^*_T(\tau;x))\,\mathrm{d}\tau - V(\tilde{x}(T;x)) + \int_{T}^{T+\zeta} q(\tilde{x}(\tau),\tilde{u}(\tau))\,\mathrm{d}\tau + V(\tilde{x}(T+\zeta))\\
&= J^*_T(x) - \int_0^{\zeta} q(x^*_T(\tau;x),u^*_T(\tau;x))\,\mathrm{d}\tau + \int_{T}^{T+\zeta} q(\tilde{x}(\tau),\tilde{u}(\tau))\,\mathrm{d}\tau + \int_{T}^{T+\zeta} \dot{V}(\tilde{x}(\tau))\,\mathrm{d}\tau\\
&\le J^*_T(x) - \int_0^{\zeta} q(x^*_T(\tau;x),u^*_T(\tau;x))\,\mathrm{d}\tau
\end{aligned}
$$

since $\int_T^{T+\zeta} q(\tilde{x}(\tau),\tilde{u}(\tau))\,\mathrm{d}\tau + \int_T^{T+\zeta}\dot{V}(\tilde{x}(\tau))\,\mathrm{d}\tau \le 0$ by (9). The result follows since the optimal cost satisfies $J^*_T(x^*_T(\zeta;x)) \le J_T(x^*_T(\zeta;x),\tilde{u}(\cdot))$. □

Note that ζ needs not be constant (uniform) for all x; it can depend on x as long as $\zeta(x)>0$, for all $x\in\Omega$.

Corollary 1. *Suppose the terminal cost function is replaced by the infinite horizon cost-to-go resulting from the application of an a priori obtained stabilizing controller that satisfies* (9) *for a sequence of control horizons* $\{\zeta_\ell\}$*, with terms defined recursively as* $\zeta_\ell = \zeta\big(x\big(\sum_{k=1}^{\ell-1}\zeta_k\big)\big)$*. Then* (11) *still holds.*

Proof. The argument follows almost inductively, based on the proof of Proposition 2:

$$
\begin{aligned}
J_T(x^*_T(\zeta_1;x),\tilde{u}(\cdot)) &= \int_{\zeta_1}^{T+\zeta_1} q(\tilde{x}(\tau),\tilde{u}(\tau))\,\mathrm{d}\tau + \int_{T+\zeta_1}^{\infty} q(x^k(\tau;x(T+\zeta_1)),k(x(\tau)))\,\mathrm{d}\tau\\
&= \int_{\zeta_1}^{T+\zeta_1} q(\tilde{x}(\tau),\tilde{u}(\tau))\,\mathrm{d}\tau + \sum_{\ell=1}^{\infty}\int_{T+\sum_{k=1}^{\ell}\zeta_k}^{T+\sum_{k=1}^{\ell+1}\zeta_k} q(x^k(\tau;x(T+\ell\zeta)),k(x(\tau)))\,\mathrm{d}\tau\\
&= J^*_T(x) - \int_0^{\zeta_1} q(x^*_T(\tau;x),u^*_T(\tau;x))\,\mathrm{d}\tau - \sum_{\ell=0}^{\infty}\int_{T+\sum_{k=1}^{\ell}\zeta_k}^{T+\sum_{k=1}^{\ell+1}\zeta_k} q(x^k(\tau;x(T+\sum_{i=1}^{\ell}\zeta_i)),k(x(\tau)))\,\mathrm{d}\tau\\
&\quad + \int_{T}^{T+\zeta_1} q(\tilde{x}(\tau),\tilde{u}(\tau))\,\mathrm{d}\tau + \sum_{\ell=1}^{\infty}\int_{T+\sum_{k=1}^{\ell}\zeta_k}^{T+\sum_{k=1}^{\ell+1}\zeta_k} q(x^k(\tau;x(T+\sum_{i=1}^{\ell}\zeta_i)),k(x(\tau)))\,\mathrm{d}\tau\\
&= J^*_T(x) - \int_0^{\zeta_1} q(x^*_T(\tau;x),u^*_T(\tau;x))\,\mathrm{d}\tau\\
&\quad + \int_{T}^{T+\zeta_1} q(\tilde{x}(\tau),\tilde{u}(\tau))\,\mathrm{d}\tau - \int_{T}^{T+\zeta_1} q(x^k(\tau;x(T)),k(x(\tau)))\,\mathrm{d}\tau\\
&= J^*_T(x) - \int_0^{\zeta_1} q(x^*_T(\tau;x),u^*_T(\tau;x))\,\mathrm{d}\tau
\end{aligned}
$$

The proof is completed by noting that the optimal cost satisfies $J^*_T(x^*_T(\zeta;x)) \le J_T(x^*_T(\zeta;x),\tilde{u}(\cdot))$. □

The following lemma is the equivalent of [3, Lemma 2.2.1] which states that if originally satisfied, condition (10) remains true under each iteration of the receding horizon scheme.

Lemma 1. *Suppose that* $x\in\Omega$*. Then* $x^*_T(T;x)\in\Omega$ *for any* $T=\sum_{i=1}^{\ell}\zeta_i$*, and for every* $\ell\ge 1$*, where* ζ_i *satisfies* (9).

Proof. The statement follows directly from the stability argument in the proof of Proposition 1, once inequality (9) is established. □

The next statement is pivotal in establishing stability for the infinite horizon. It is the equivalent of [3, Prop 2.2.2], stated and proved for a function that satisfies the requirements of Proposition 1 in the role of the terminal cost function in (6).

Proposition 3. *Let $T > 0$, be such that $x^*_T(T;x) \in \Omega$ for all $x \in \Gamma$, where Γ is the path connected component of Ω in which $J^*_T(x)$ is finite.*[2] *Let $x_0 \in \Gamma$ and consider a trajectory $(x_{\mathrm{rh}}(t), u_{\mathrm{rh}}(t))$, $t \geq 0$, resulting from the use of a receding horizon strategy $\mathscr{RH}(T, \{\zeta_k\})$ with $\delta_k > 0$, $\sum_{k=0^\ell} \zeta_k \to \infty$ as $\ell \to \infty$. Then,*

$$J_\infty(x_0, u_{\mathrm{rh}}(\cdot)) \leq J^*_T(x_0) \quad .$$

Proof. The receding horizon strategy defines a sequence of points $\{x_k\}_{k=0}^\infty$ according to $x_{k+1} = x^*_T(\zeta_k, x_k)$, starting with x_0 so that $x_k = x(\sum_{i=1}^k \zeta_i)$. By the principle of optimality, the cost of flowing from x_k to x_{k+1} is given by

$$\int_{\sum_{i=1}^k \zeta_i}^{\sum_{i=1}^{k+1} \zeta_i} q\big(x_{\mathrm{rh}}(\tau), u_{\mathrm{rh}}(\tau)\big)\, d\tau = J^*_T(x_k) - J^*_{T-\zeta_{k+1}}(x_{k+1}) \geq 0 \quad .$$

The total cost of this strategy is

$$J_\infty(x_0, u_{\mathrm{rh}}(\cdot)) = J^*_T(x_0) + \sum_{k=1}^\infty \{J^*_T(x_k) - J^*_{T-\zeta_{k+1}}(x_k)\} \leq J^*_T(x_0) \quad ,$$

where the final inequality follows from the fact that $J^*_T(x_k) \leq J^*_{T-\zeta}(x_k)$, for all $\zeta \geq 0$ satisfying (9) and $k \geq 0$. Indeed,

$$\begin{aligned}
J^*_T(x) &= \int_0^{T-\zeta} q(x^*_T(\tau;x), u^*(\tau;x))\,\mathrm{d}\tau + \int_{T-\zeta}^{T} q(x^*_T(\tau;x), u^*(\tau;x))\,\mathrm{d}\tau + V(x^*_T(T;x)) \\
&\leq \int_0^{T-\zeta} q(x^*_T(\tau;x), u^*(\tau;x))\,\mathrm{d}\tau + \int_{T-\zeta}^{T} q\big(x^k_T(\tau;x), u^k(\tau;x)\big)\,\mathrm{d}\tau + V(x^*_T(T;x)) \\
&\leq \int_0^{T-\zeta} q(x^*_T(\tau;x), u^*(\tau;x))\,\mathrm{d}\tau - \int_{T-\zeta}^{T} \dot{V}\big(x^*_T(\tau;x), u^*(\tau;x)\big)\,\mathrm{d}\tau + V(x^*_T(T;x)) \\
&= \int_0^{T-\zeta} q(x^*_T(\tau;x), u^*(\tau;x))\,\mathrm{d}\tau + V(x^*_T(T-\zeta;x)) = J^*_{T-\zeta}(x) \quad .
\end{aligned}$$

□

The sequence of results in this section culminates to the following (exponential) stability statement, the analog of [3, Thm 2.2.2], established similarly thanks to Assumption 1, which although reasonable, may be somewhat restrictive. The result is existential, since estimating the constants in Assumption 1 may not be practical.

Theorem 1. *Let $T > 0$ and consider the use of a receding horizon scheme $\mathscr{RH}(T, \{\zeta_k\})$ with each $\zeta_k \in (0, T]$ and $\sum_{j=0}^k \zeta_j \to \infty$ as $k \to \infty$. Then, for each $x_0 \in \Gamma$, the resulting trajectory converges to the origin exponentially fast.*

Proof. Define the continuous function

$$W(t; x_0, u_{\mathrm{rh}}(\cdot)) \triangleq \int_t^\infty q(x_{\mathrm{rh}}(\tau), u_{\mathrm{rh}}(\tau))\,\mathrm{d}\tau \quad .$$

By Proposition 3, $W(t; x_0, u_{\mathrm{rh}}(\cdot)) \leq J^*_T(x_{\mathrm{rh}}(t))$, and since $\mathscr{RH}(T, \{\delta_k\})$ is suboptimal, $J^*_\infty(x_{\mathrm{rh}}(t)) \leq W(t; x_0, u_{\mathrm{rh}}(\cdot))$. Now,

$$\frac{\partial}{\partial t} W(t; x_0, u_{\mathrm{rh}}(\cdot)) = -q(x_{\mathrm{rh}}(t), u_{\mathrm{rh}}(t)) \leq -c_q \|x_{\mathrm{rh}}(t)\|^2 \leq -\frac{c_q}{c_T} J^*_T(x_{\mathrm{rh}}(t)) \leq -\frac{c_q}{c_T} W(t; x_0, u_{\mathrm{rh}}(\cdot)) \quad .$$

[2] Since $V(x)$ is finite in Ω by definition, this condition depends on how $q(x, u)$ is defined.

From the Comparison Lemma it follows that

$$W(t;x_0,u_{\mathrm{rh}}(\cdot)) \leq e^{-\frac{c_q}{c_T}\,t}\,W(0;x_0,u_{\mathrm{rh}}(\cdot))$$
$$\Rightarrow m_\infty \|x_{\mathrm{rh}}(t)\|^2 \leq J_\infty^*(x_{\mathrm{rh}}(t)) \leq W(t;x_0,u_{\mathrm{rh}}(\cdot)) \leq e^{-\frac{c_q}{c_T}\,t}\,J_T^*(x_0) \leq c_T\; e^{-\frac{c_q}{c_T}\,t}\,\|x_0\|^2 \;\;,$$

which implies that $\|x_{\mathrm{rh}}(t)\|$ converges exponentially to the zero with a rate lower bounded by $\frac{c_q}{c_T}$. □

Thus, Proposition 3 and Assumption 1 establish stability for the receding horizon scheme. Proposition 3 is based on Proposition 1, which relaxes sign definiteness for the derivative of the terminal cost in exchange for constraints on switching times. Among the conditions placed for stability, the ones in Assumption 1, although reasonable, are probably the most difficult to formally verify, whereas the conditions of Proposition 1 will be enforced by construction in the section that follows.

In conclusion, it is noted that although the approach described in this section is motivated by the integration of navigation functions in the receding horizon framework, the results of Section 2 suggest that the terminal cost estimate *does not need to be a navigation function*. Rather, any positive definite function satisfying the requirements of Proposition 1 can qualify as a terminal cost. Thus, Proposition 1 practically allows the use of a more general class of potential functions, as long as the dynamics of the system can ensure the satisfaction of (4) and (5).

4 Randomized Model Predictive Control

While for linear control systems, the structure of the model admits closed form solutions and available solvers suffice both in terms of computational speed as well as complexity, this might not necessarily be the case for nonlinear systems, where no closed form solution is known and the search for the optimum may rely heavily on initial guess, model size, complexity, and stability properties, as well as the presence of inequality state constraints [9].

This section offers an alternative approach to finite-time optimal control design, using randomized methods. Instead of numerically computing an optimal control law $u^*(\cdot)$ that minimizes the functional (6), the proposed method computes *probable near minimizers* of the finite-horizon cost, at chosen levels of confidence. This method is based on the VC-theory of statistical learning, specifically the approach to distribution-free learning along the lines of [2] that uses Chernoff bounds to determine the size of an appropriately large sample. The proposed approach thus builds the optimal control strategy on a completely different computational framework that is based on sampling, without relying on specialized nonlinear optimal control design numerical tools. This approach is motivated by the fact that the stability analysis of the preceding section does not necessarily require the feedback controller $k(x)$ to be truly optimal in order to guarantee stability; rather, any feedback controller $k(x)$ that provides a sufficient amount of decrease for the finite horizon cost is adequate [3]. For completeness, the following proposition is stated:

Proposition 4. *For a fixed $T > 0$, and for $\zeta_i > 0$ for every $i \in \mathbb{N}_+$, let $\{u_i(x)\}_{i=1}^\infty$ be a sequence of control laws satisfying* (9) *for which*

$$J_T\big(x_{i+1},u_{i+1}(\cdot)\big) < J_{T-\zeta_i}\big(x_{i+1},u_i(\cdot+\zeta_i)\big) \;\;, \tag{12}$$

for all $i \in \mathbb{N}_+$. Then $x^{u_i}(t) \to 0$ as $t \to \infty$.

Proof. Condition (12) essentially states that switching from u_i to u_{i+1} at x_{i+1} improves the value of the finite horizon cost at that state (note that for $t = 0$, $u_i(t)$ needs to give the control input at x_i). Define the sequence of finite horizon costs $\{c_i\}_{i=1}^\infty$, with $c_i \triangleq J_T\big(x_i,u_i(\cdot)\big)$, and observe that for $T = \zeta_i$

$$\begin{aligned}
c_i - c_{i+1} &= J_{\zeta_i}\big(x_i,u_i(\cdot)\big) - J_{\zeta_i}\big(x_{i+1},u_{i+1}(\cdot)\big) > J_{\zeta_i}\big(x_i,u_i(\cdot)\big) - J_{\zeta_i-\zeta_i}\big(x_{i+1},u_i(\cdot+\zeta_i)\big)\\
&= \int_0^{\zeta_i} q\big(x^{u_i}(\tau;x_i),u_i(\cdot)\big)\,\mathrm{d}\tau + V\big(x^{u_i}(\zeta_i;x_i)\big) - V\big(x^{u_{i+1}}(0;x_{i+1})\big)\\
&= \int_0^{\zeta_i} q\big(x^{u_i}(\tau;x_i),u_i(\cdot)\big)\,\mathrm{d}\tau + V(x_{i+1}) - V(x_{i+1}) = \int_0^{\zeta_i} q\big(x^{u_i}(\tau;x_i),u_i(\cdot)\big)\,\mathrm{d}\tau \;\;.
\end{aligned}$$

Condition (9) then ensures that $c_i - c_{i+1} > \gamma(\|x_i\|)$ suggesting that $\{c_i\}$ is monotonically decreasing. In addition, it is nonnegative by definition, and therefore converges. If that limit is assumed to be bounded away from zero (implying that $q(x,u) > 0$), then $c_i - c_{i+1} \to 0$, while $\gamma(\|x_i\|) \to c > 0$, which is a contradiction. Therefore one has to have $c_i \to 0 \Rightarrow q(x,u) \to 0 \Rightarrow x \to 0$. □

The working assumption made here is that each feedback control law $u = k(x)$ can be parameterized in terms of a set of control parameters $\eta \in H$. Suppose that $\mathbb{P}$ is a probability measure on H, and denote the underlying probability distribution (prior) on η $P(\eta)$. With the given parameterization, the cost of applying input $u = k(\eta;t,x)$ from x over the time horizon T is rewritten as

$$J_T(x,\eta) = \int_0^T q(x^\eta(\tau;x), u^\eta(\tau))\, \mathrm{d}\tau + V(x^\eta(T;x)) \quad ,$$

where $u^\eta(\cdot) \triangleq k(\eta;t,x)$ and $x^\eta(\tau;x)$ is the trajectory resulting from the application of control law u^η to the system starting from x at $\tau = 0$. Thus, given a state x, the finite horizon cost becomes a function of the parameter (decision) vector η.

If the parameterization chosen is complete, in the sense that all possible feedback functions $u = k(x)$ can be expressed as a function of a (possibly infinite dimensional) decision vector η, then the optimal finite horizon cost from state x can be written

$$J_T^*(x) = \inf_{\eta \in H} J_T(\eta)$$

and the optimal decision vector η^* is such that $J_T(x,\eta) = J_T^*(x)$.

The optimization of functional $J_T(x, u(\eta;t))$ is attempted by means of randomized algorithms. Take N_s independent and identically distributed random samples μ_i, $i = 1,\dots,N_s$ from H according to a probability distribution $P(\eta)$. Then $J_T^\circ \triangleq \min_i J(\eta_i)$ can approximate J_T^* [12]:

Definition 2 (Probable Near Minimum [2])**.** *Given $J_T(\eta)$, $\delta \in (0,1)$, $\alpha \in (0,1)$, a number $J_T^\circ \in \mathbb{R}$ is said to be a probable near minimum of $J_T(\eta)$ to level α and confidence $1-\delta$ if there exists a set $\tilde{H} \subseteq H$ measuring $\mathbb{P}\{\tilde{H}\} \leq \alpha$ such that*

$$\mathbb{P}\Big\{ \inf_H J_T(\eta) \leq J_T^\circ \leq \inf_{H \setminus \tilde{H}} J_T(\eta) \Big\} \geq 1 - \delta \quad . \tag{13}$$

Parameter α (level) quantifies the set of solutions that may not be represented in the samples. For small α, the probability of finding a solution better than J_T° is small. Confidence $1-\delta$ expresses the probability that the solution is not worse than the actual minimum in the set that is represented in the samples. The desired level and confidence is achieved by adjusting the number of samples N_s:

Lemma 2 (Number of samples [12, 2])**.** *The number of samples N_s that guarantees that J_T° is a probable near minimum of $J_T(\eta)$ to level α and confidence $1-\delta$ satisfies*

$$N_s \geq \frac{\ln(1/\delta)}{\ln(1/(1-\alpha))} \quad . \tag{14}$$

Lemma 2 provides the basis for the following algorithm, which computes a probable minimizer of the finite horizon cost through sampling. The algorithm is a version of the distribution-free randomized optimization algorithms of [12, 2], modified to account for the additional stability constraints on the parameterized inputs. In this particular implementation, the algorithm is simplified by setting a constant control horizon ζ.

Algorithm 1 Randomized optimization

Require: α, δ, $J_T(\eta)$, $P(\eta)$, H, T, ζ, $\gamma(\cdot)$, $\theta(\cdot)$, x, $x^\eta(t,x)$, $V(x)$, $\dot{V}(\eta;x)$.
Ensure: J_T° and η°.

1: $N_s \leftarrow \lceil \frac{\ln(1/\delta)}{\ln(1/(1-\alpha))} \rceil$.
2: $\tilde{\eta} \leftarrow \emptyset$
3: **repeat**
4: Generate $N_s - |\tilde{\eta}|$ i.i.d. samples η_i from H according to $P(\eta)$
5: $\tilde{\eta}^+ \leftarrow \{\eta_i \mid V(x^\eta(\zeta,x)) - V(x) \leq \gamma(\|x\|) \wedge \max_{t\in[0,\zeta]} |\dot{V}(\eta_i; x^\eta(t,x))| \leq \frac{\theta(\|x\|)}{\zeta}\}$
6: $\tilde{\eta} \leftarrow \tilde{\eta} \cup \tilde{\eta}^+$
7: **until** $|\tilde{\eta}| = N_s$
8: $J_T^\circ \leftarrow \min_{\tilde{\eta}} J_T(\eta_i)$
9: $\eta^\circ \leftarrow \arg\min_{\tilde{\eta}} J_T(\eta_i)$

Minimizing the functional $J_T(x,u)$ within the receding horizon scheme is therefore achieved by parameterizing the set of feedback control laws, sampling the subset of the parameter space to obtain a sufficient number of samples consistent with the stability constraints of Proposition 1 (step 5), and evaluating an admissible probable near minimizer $u^{\circ}(\eta;x)$.

The algorithm will halt at step 3 if the stability constraints cannot be satisfied with a sufficient number of sample parameters. Alternatively, one can set an upper bound on the total number of samples drawn, and have the algorithm report its solution with a confidence that is accordingly reduced. Increasing δ reduces N_s and consequently the number of iterations in the loop of step 3. This is the main mechanism the algorithm provides to trade-off performance for computational speed.

5 Numerical Implementation and Testing

The proposed control synthesis methodology is implemented numerically, for a spherical, two-dimensional workspace including a single rectangular obstacle. The boundary of the workspace is defined by that of a ball, centered at configuration $(-3,3)^T$ and having a radius of 5: $\Omega \triangleq \{q \in \mathbb{R}^2 \mid \|q-(-3,3)^T\| \leq 5\}$. The destination configuration is set at $q_d = (-4,3)^T$. (To conform with the assumptions of Proposition 1, one can consider a change of coordinates placing the new origin at $(-4,3)^T$.) While moving within this workspace, it is assumed that the robot needs to communicate wirelessly with a radio base station, located at position $(-4,5)$, and the degradation of the radio signal due to shadowing by the obstacle needs to be taken into account during navigation. A navigation function $\varphi(q)$ is constructed (Fig. 2).

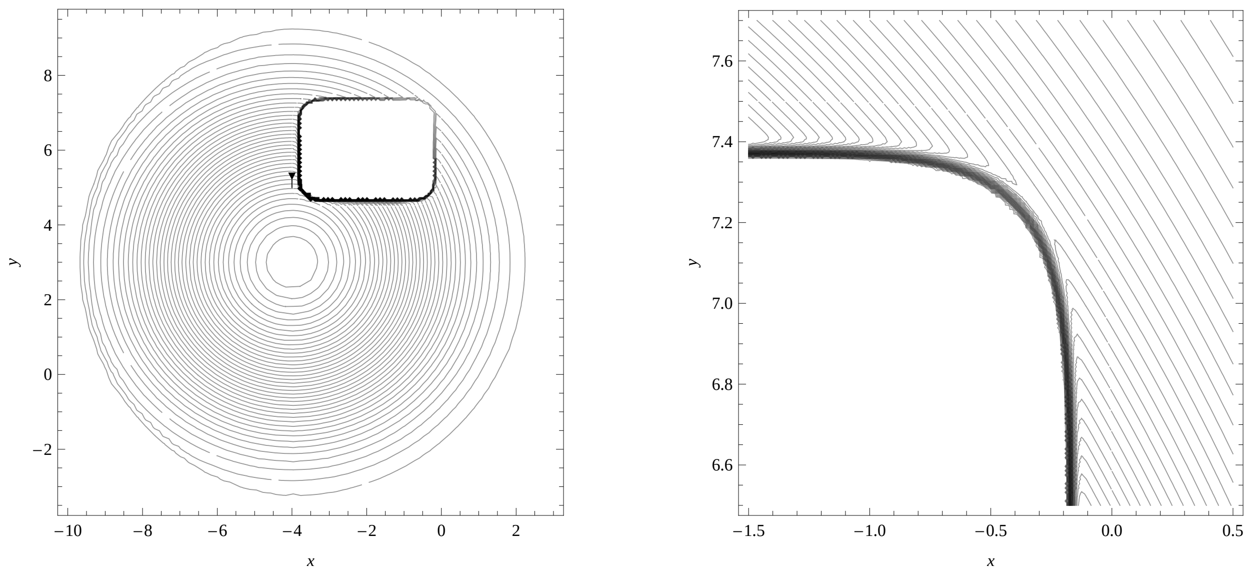

(a) Contour plot of the navigation function on a spherical workspace with a rectangular obstacle. The destination configuration is slightly off-center. The position of the radio base station is marked close to the bottom left hand corner of the obstacle.

(b) Detail of the potential field behind the obstacle. There is a saddle point close to the obstacle boundary, but no local minimum in the neighborhood.

Figure 2: The navigation function used in the numerical implementation example.

The system tested is one with simple dynamics in the form of a single integrator $\dot{q} = u$, where $q = (x,y)^T$ and $u \in \mathbb{R}^2$. To parameterize the input, we use a basis of Legendre polynomials with five elements $\hat{p}_0(t),\ldots,\hat{p}_4(t)$ (Fig. 3(a)). The basis polynomials are scaled in time to have a domain $t \in [0,T]$ where $T = 10$ seconds is set to be the prediction horizon (Fig. 3(b)). A uniform control horizon of $\zeta = 5$ seconds is selected. Relatively large control horizons are possible since the system does not run entirely open-loop during this time as indicated in (15), although (4) introduces a trade-off between the selected control horizon and the range of admissible inputs through (16). On the other hand, ζ should be comparable to the time scale of the basis functions, otherwise the selection of the admissible inputs becomes very conservative, making the method inefficient.

A linear combination of the scaled Legendre polynomials define a time-varying perturbation on the direction of the potential field. A control input is thus parameterized through a set of five coefficients $\eta = (\eta^{(0)}, \ldots, \eta^{(4)})$ in polar coordinates, and the input sample includes the nominal negated gradient:

$$\begin{aligned} u_0(q) &= -\nabla_q \varphi(q) \\ u_i(t,q) &= \left(\|\nabla \varphi(q)\|, \arg\left(-\nabla\varphi(q) + \frac{\pi}{2} \sum_{i=0}^{4} \eta^{(i)}\, p_i(t) \right) \right) \qquad i = 1, \ldots, N_s \enspace , \end{aligned} \tag{15}$$

where $\arg(\cdot)$ denotes the angle of its vector argument.

For the randomized algorithm a confidence $\delta = 0.1$ and a level $\alpha = 0.1$ is selected. This yields a sample size of $N_s = 22$. To generate a random sample for the control inputs in the form (15), a collection of N_s sets of parameters η are uniformly sampled from the set $[-\frac{3}{5}, \frac{3}{5}]$, the size of which can either be determined empirically or analytically in order to keep the generated perturbation on the nominal direction within certain bounds with given probability. (For example, sampling from $[-\frac{1}{5}, \frac{1}{5}]$ keeps *all* angle perturbations within the $[-\frac{\pi}{2}, \frac{\pi}{2}]$ interval.)

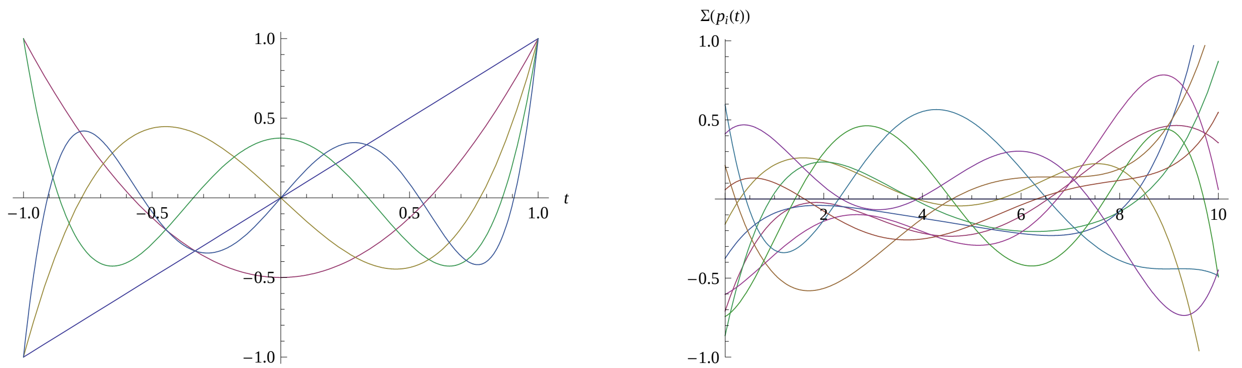

(a) Graphs of the Legendre polynomials used as basis for the parameterized family of input functions.

(b) Ten different parameterizations using the Legendre polynomial basis, scaled to take values in the (0,T) interval, where T is the prediction horizon.

Figure 3: A Legendre polynomial basis is chosen to parameterize the control inputs. Each input candidate is a perturbation over the (nominal) direction of the potential field as indicated in (15).

The system is initialized at location $q_0 = (0.3, 6)^T$, which corresponds to a configuration where the rectangular obstacle obscures the line of sight both to the destination configuration and the radio base station (Fig. 2). At this configuration, the navigation function assumes the value of $\varphi(q_0) = 0$. A class-$\mathscr{K}_\infty$ function θ lower bounds $\varphi(q)$ within Ω. Condition (4) of Proposition 1 is enforced by regulating the input of (15) for $t \in (\tau_i, \tau_{i+1}]$ as follows

$$u(t,q) = \left(\left| V\left(t, q; q(\tau_i)\right) \right|, \arg\left(-\nabla\varphi(q) + \frac{\pi}{2} \sum_{i=0}^{4} \eta^{(i)}\, p_i(t) \right) \right) \enspace ,$$

where

$$V(t,q;q(\tau_i)) \triangleq \min\left\{ \frac{\theta(\|q(\tau_i) - q_d\|)}{\zeta}, \operatorname{sign}\left(\left| \sum_{i=0}^{4} \eta^{(i)}\, p_i(t) \right| - \frac{\pi}{2} \right) \|\nabla\varphi(q)\| \right\} \enspace , \tag{16}$$

and ζ denotes the control horizon. The construction in (16) compares the magnitude of the gradient $\|\nabla\varphi\|$ to the class-$\mathscr{K}_\infty$ function $\frac{\theta(\cdot)}{\zeta}$, and caps the magnitude of the input vector to the latter if the input sends the system to higher levels of $\varphi(q)$. To eliminate the possibility of collision (reaching the boundary of the free space), it suffices to have

$$\theta(\|q_0 - q_d\|) < 1 - \varphi(q_0) \enspace , \tag{17}$$

and for the particular choice of initial conditions, an appropriate selection for $\theta(\cdot)$ that satisfies both the requirements of Proposition 1 and (17) is $\theta(\|q - q_d\|) = 0.004\|q - q_d\|^2$. For the examples presented a class $\mathscr{K}_\infty$ function $\gamma(r) = 10^{-6} r^2$ is used.

The finite-horizon cost functional is defined as

$$J_T(q,u) = \int_0^T \{c_1 \|u(\tau)\|^2 + c_2 \|x(\tau)\|^2 + c_3\, C\big(x(\tau)\big)\}\, \mathrm{d}\tau + \varphi(x) \quad , \tag{18}$$

where c_1, c_2, and c_3 are positive constants, and $C(x)$ quantifies the cost of communication between the robot and the radio base station. This cost is expressed as a integral of a function $\beta(x)$ that captures the obstacle shape (Fig. 4) over the line that connects the robot to the base station,

$$C(x) = \int_0^1 \beta\big((1-\chi)x + \chi x_r\big) \mathrm{d}\chi \quad ,$$

where x_r denotes the position of the base station. As shown in Fig. 4, the function is constructed in such a way so that if there is line of sight, the shadowing cost is zero. For this implementation, $c_1 = c_2 = 10^{-5}$, and $c_3 = 0.1$.

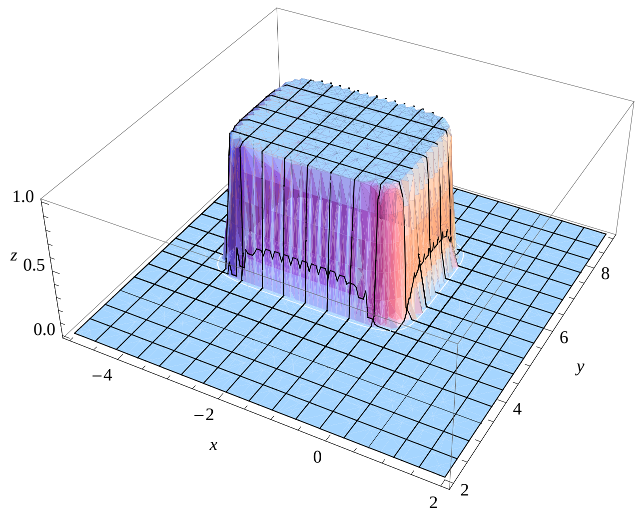

Figure 4: Graph of function β that expresses the shape of the obstacle which introduces shadowing effects on the radio signal.

Figure 5 illustrates an output of the randomized receding horizon scheme in the form of a path from the chosen initial condition. Figure 5(a) depicts a partial path from initial configuration to goal, with the intention of comparing mid-way the performance of the randomized receding horizon scheme to that of the standard steepest descent in terms of convergence speed. Identifying the initial and goal configurations in Fig. 5(a), as $(0.3, 6)$ and $(-4, 3)$, respectively, the figure shows that at the end of the fifth step, the randomized receding horizon scheme brings the robot closer to the destination than steepest descent. Toward the end of the simulation time frame, when a line of sight with the radio base station has been established and no obstacle is in the way, the finite-time probable near optimal inputs coincide with the negated potential function gradient.

Figure 6 shows the evolution of the incremental and terminal cost with time, along the probable near optimal path drawn by the solid curve in Fig. 5(a). It is shown that the terminal cost is allowed to increase temporarily if the candidate solution shows promise in a longer horizon. This issue is explored in more detail in Fig. 7 and in the paragraph that follows.

A local minimum is artificially introduced in function $\varphi(q)$ used (18), so that it is no more a navigation function. The local minimum is created at location $(-1.5, 4)$. In the simulation run that uses the modified potential function, the objective is to test the ability of the randomized receding horizon scheme to work around local minima, based on the relaxed conditions on the evolution of the terminal cost imposed in Proposition 1. The results are presented in Fig. 7, where it is shown that the strategy is capable of finding an alternative route to avoid the minimum.

6 Conclusions and Outlook

Lyapunov-based conditions for asymptotic stability of switching nonlinear systems can be relaxed, so that the Lyapunov-like function need not have a negative semi-definite derivative, as long as one can guarantee that the function will have

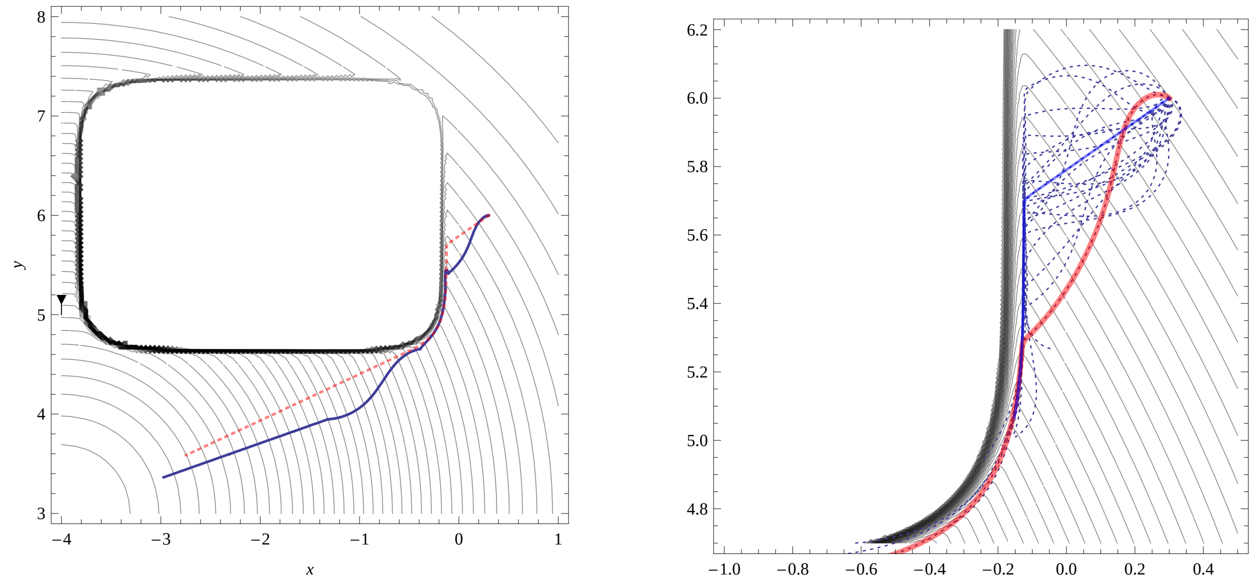

(a) The solid curve represents the path produced by the randomized receding horizon scheme, while the dashed curve shows the path corresponding to steepest descent. Note how the receding horizon path curves to the bottom-left to reduce the communication shadowing effect from the obstacle.

(b) An example of candidate paths resulting from different input samples. Different paths are shown in dashed curves. The steepest decent path is shown with the solid curve and for the prediction horizon selected, terminates close to $(-0.175, 5.1)$. The probable near optimal one is highlighted with the thick solid line, and continues beyond the selected frame.

Figure 5: Randomized receding horizon robot navigation paths. Even though the receding horizon path is longer, and despite penalizing motion speed, the receding horizon path still reaches the destination faster than the steepest descent. The trajectories in 5(b) are from a different simulation run compared to the ones in 5(a). In 5(a), the robot briefly increases its potential along its probable near optimal trajectory before it reduces it again, whereas in 5(b) the probable near optimal trajectory found monotonically reduces the potential function.

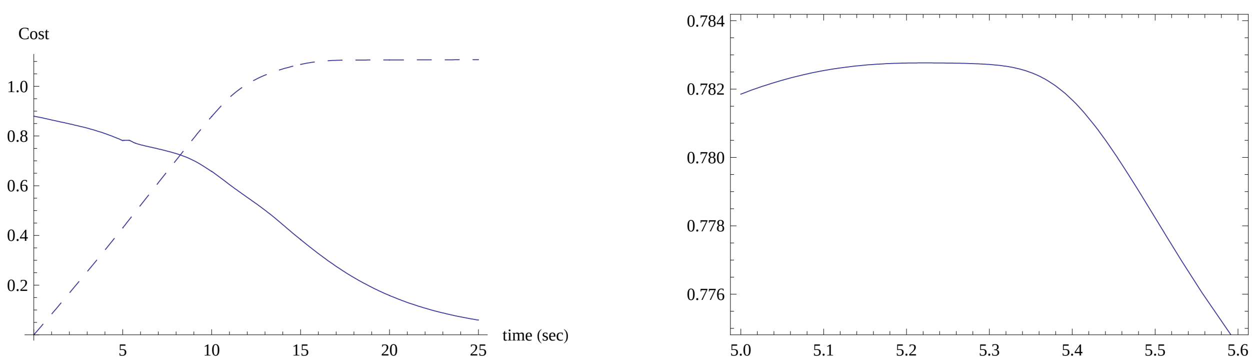

(a) Evolution of terminal cost (shown in solid curve) and incremental cost (shown in dashed curve) along a solution of the receding horizon scheme over time. The simulation period covers five steps, each with a control horizon of five seconds. Note the brief increase on the terminal cost close to the fifth simulation second, shown in detail in Fig. 6(b)

(b) The terminal cost, expressed by the value of the navigation function, temporarily increases along the probable near optimal trajectory, consistently with the requirements of Proposition 1. This brief increase may be a result of a sample input that yielded significantly better performance compared to other candidate solutions in the sample set.

Figure 6: Behavior of the incremental and terminal costs defined in (18), along the randomized receding horizon trajectory shown in Fig. 5(a).

significantly decreased between switching times. In some sense, this may be equivalent to looking at the sampled data system at switching times, and applying a Lyapunov-type argument for the sampled data system. This approach overcomes a technical limitation to using a navigation function as a control Lyapunov function in nonlinear systems with

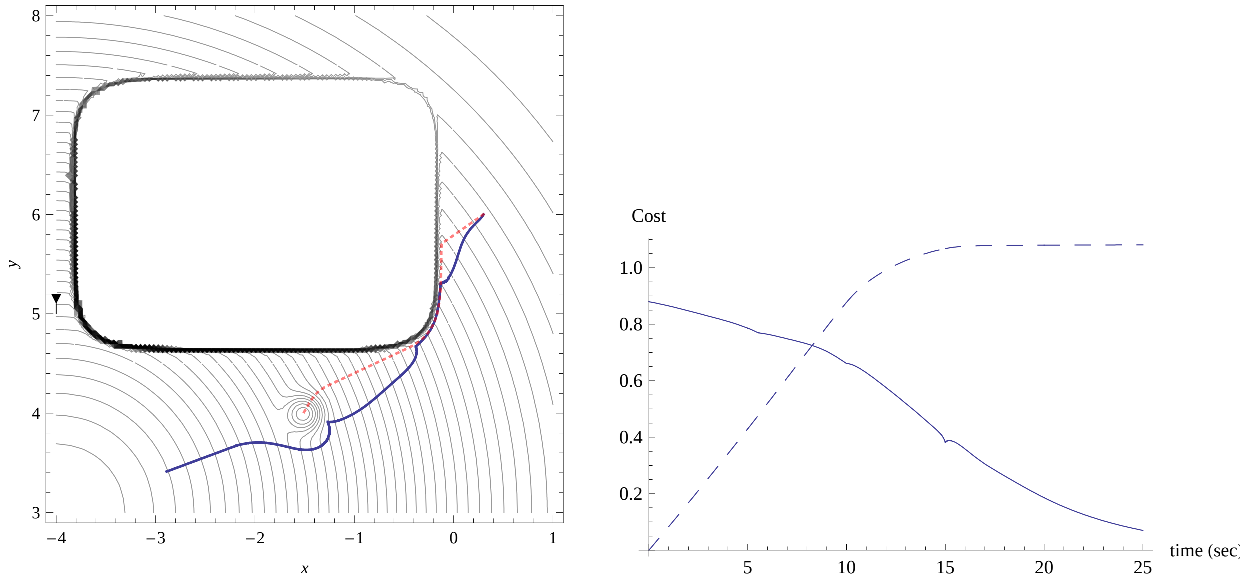

(a) Oups! pitfall. In this simulation run of four steps, there is a local minimum at which the steepest descent solution shown with the dashed curve gets trapped. The randomized receding horizon solution circumvents the minimum and continues toward the goal, bending to the right trying to simultaneously decrease the shadowing effect of the obstacle.

(b) Terminal and incremental cost evolution along the randomized receding horizon path of Fig. 7(a). At the beginning of the fourth simulation step, the algorithm decides to increase the terminal cost temporarily, so that the local minimum is circumvented.

Figure 7: The relaxed conditions of a receding horizon scheme based on Proposition 1 allow the use of terminal costs that cannot qualify as control Lyapunov functions, even for cases of the later with negative semidefinite time derivatives. This potential function has a local minimum at $(-1.5, 4)$, yet the robot under the randomized receding horizon strategy circumvents it on its way to destination.

switching dynamics, and it also allows this class of functions to be used as terminal costs in receding horizon schemes. When coupled with a receding horizon scheme, these terminal cost functions do not need to be navigation functions, and it is suggested that a more general class of positive definite potential functions can qualify for use in a receding horizon mobile robot navigation control strategy. Finally, the finite horizon optimal control design that is part of the receding horizon scheme can be implemented using a class of randomized algorithms based on Chernoff bounds. The use of this type of statistical learning algorithms can demonstrate how the combination of a receding horizon scheme with a potential function on the robot's workspace can be successfully used to circumvent the problem of local minima associated with typical potential functions, which are not constructed within the navigation function framework.

One of the remaining challenges ongoing research is targeting is handling uncertainty, both at the level of the system dynamics, and at that of environment. It is anticipated that the requirements of statistical learning theory associated with the P-dimension of the objective function will present difficulties related to inability to ensure uniform convergence of an empirical mean, and may necessitate the reformulation of Algorithm 1.

Acknowledgements

This work is supported by ARL through BAE Systems as part of MAST CTA #W911NF-08-2-0004.

References

[1] J. Piovesan and H. Tanner, "Randomized model predictive control for robot navigation," Submitted to IEEE ICRA, 2009.

[2] M. Vidyasagar, “Randomized algorithms for robust controller synthesis using statistical learning theory,” *Automatica, Elsevier Science Ltd.*, vol. 37, pp. 1515–1528, 2001.

[3] A. Jadbabaie, “Receding horizon control of nonlinear systems: A control lyapunov function approach,” Ph.D. dissertation, California Institute of Technnology, Pasadena, California, October 2000.

[4] A. Jadbabaie, J. Yu, and J. Hauser, “Stabilizing receding horizon control of nonlinear systems: A control Lyapunov function approach,” in *Proceedings of the American Control Conference*, San Diego, CA, USA, 1999, pp. 1535–1539.

[5] J. Primbs, V. Nevistić, and J. Doyle, “Nonlinear optimal control: A control Lyapunov function and receding horizon perspective,” *Asian Journal of Control*, vol. 1, pp. 14–24, 1999.

[6] G. Grimm, M. J. Messina, S. E. Tuna, and A. R. Teel, “Model predictive control: For want of a local control lyapunov function, all is not lost,” *IEEE Transactions on Automatic Control*, vol. 50, no. 5, pp. 546–558, May 2005.

[7] S. S. Keerthi and E. G. Gilbert, “Optimal infinite-horizon feedback laws for a general calss of constrained discrete time systems: Stability and moving-horizon approximations,” *Journal of Optimization Theory and Applications*, vol. 57, no. 2, pp. 265–293, 1988.

[8] A. Jadbabaie and J. Hauser, “On the stability of unconstrained receding horizon control with a general terminal cost,” in *Proceedings of the Conference on Decision and Control*, 2001, pp. 4826–4831.

[9] A. Schwartz, “Theory and implementation of numerical methods based on runge-kutta integration for optimal control problems,” Ph.D. dissertation, University of California, Berkeley, 1996.

[10] M. B. Milam, K. Mushambi, and R. M. Murray, “A new computational approach to real-time trajectory generation for constrained mechanical systems,” in *Proceedings of the Conference on Decision and Control*, 2000, pp. 845–851.

[11] M. Vidyasagar, *A Theory of Learning and Generalization*, ser. Communications and Control Engineering. London, UK: Springer, 1997.

[12] R. Tempo, G. Calafiore, and F. Dabbene, *Randomized Algorithms for Analysis and Control of Uncertain Systems*, ser. Communications and Control Engineering. London, UK: Springer, 2003.

[13] A. Nemirovski, “Several NP-hard problems arising in robust stability analysis,” *Mathematics of Control, Signals, and Systems*, vol. 6, no. 2, pp. 99–105, June 1993.

[14] G. Calafiore, F. Dabbene, and R. Tempo, “A survey of randomized algorithms for control synthesis and performance verification,” *Journal of Complexity*, vol. 23, no. 3, pp. 301–316, June 2007.

[15] V. Koltchinskii, C. Abdallah, M. Ariola, P. Dorato, and D. Panchenko, “Improved sample complexity estimates for statistical learning control of uncertain systems,” *IEEE Transactions on Automatic Control*, vol. 45, no. 12, pp. 2383–2388, December 2000.

[16] ——, “Statistical learning control of uncertain systems: It is better than it seems,” ECE - University of New Mexico, Albuquerque, NM, USA, Tech. Rep., 2000.

[17] G. Calafiore and B. T. Polyak, “Stochastic algorithms for exact and approximate feasibility of robust lmis,” *IEEE Transactions on Automatic Control*, vol. 46, no. 11, pp. 1755–1759, November 2001.

[18] Y. Fujisaki and Y. Oishi, “Guaranteed cost regulator design: A probabilistic solution and a randomized algorithm,” *Automatica*, vol. 43, no. 2, pp. 317–324, February 2007.

[19] H. Ishii, T. Basar, and R. Tempo, “Randomized algorithms for synthesis of switching rules for multimodal systems,” *IEEE Transactions on Automatic Control*, vol. 50, no. 6, pp. 754–767, June 2005.

[20] J. F. Canny, *The Complexity of Robot Motion Planning*. MIT Press, 1998.

[21] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *The International Journal of Robotics Research*, vol. 5, no. 1, pp. 396–404, Spring 1986.

[22] E. U. Acar, H. Choset, A. A. Rizzi, P. N. Atkar, and D. Hull, "Morse decompositions for coverage tasks," *The International Journal of Robotics Research*, vol. 21, no. 4, pp. 331–344, April 2002.

[23] L. Kavraki, P. Švestka, J.-C. Latombe, and M. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.

[24] J. Barraquand, B. Langlois, and J.-C. Latombe, "Numerical potential fields techniques for robot path planning," *IEEE Transactions on Systems, Man and Cybernetics*, vol. 22, pp. 224–241, 1992.

[25] B. Krogh, "A generalized potential field approach to obstacle avoidance control," in *Proc. of SME Conference on Robotics Research*, Bethlehem, PA, 1984.

[26] A. Bemporad, A. D. Luca, and G. Oriolo, "Local incremental planning for car-like robot navigating among obstacles," in *Proc. of IEEE Conference on Robotics and Automation*, Minneapolis, MN, April 1996, pp. 1205–1211.

[27] J. Kim and P. Khosla, "Real-time obstacle avoidance using harmonic potential functions," *IEEE Transactions on Robotics and Automation*, vol. 8, no. 3, pp. 338–349, 1992.

[28] E. Rimon and D. Koditschek, "Exact robot navigation using artificial potential functions," *IEEE Transactions on Robotics and Automation*, vol. 8, no. 5, pp. 501–518, October 1992.

[29] G. P. Roussos, G. Chaloulos, K. J. Kyriakopoulos, and J. Lygeros, "Control of multiple non-holonomic air vehicles under wind uncertainty using model predictive control and decentralized navigation functions," in *Proceedings of the Conference on Decision and Control*, 2008, pp. 1225–1230.

[30] H. Choset, K. M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. Kavraki, and S. Thrun, *Principles of Robot Motion: Theory, Algorithms and Implementations*. MIT Press, 2005.

[31] D. E. Koditschek and E. Rimon, "Robot navigation functions on manifolds with boundary," *Advances in Applied Mathematics*, vol. 11, pp. 412–442, 1990.

[32] Y. Lin and E. D. Sontag, "Control-lyapunov universal formulas for restricted inputs," *Control Theory and Advanced Technology*, vol. 10, no. 4, pp. 1–22, 1995.

[33] D. Liberzon, *Switching in Systems and control*. Boston, MA, USA: Birkhauser, 2003.

[34] R. DeCarlo, M. Branicky, S. Pettersson, and B. Lennartson, "Perspectives and results on the stability of hybrid systems," *Proceedings of the IEEE*, vol. 88, no. 7, pp. 1069–1082, Jul 2000.

[35] J. Hespanha, "Uniform stability of switched linear systems: Extensions of LaSalle's invariance principle," *IEEE Trans. on Automatic Control*, vol. 49, no. 4, pp. 470–482, Apr. 2004.

[36] Y. S. Ledyaev and E. D. Sontag, "A lyapunov characterization of robust stabilization," *Nonlinear Analysis*, vol. 37, pp. 813–840, 1999.