

Randomized Receding Horizon Navigation

Herbert G. Tanner, *Senior Member, IEEE* and Jorge L. Piovesan *Member, IEEE*

Abstract—The note combines (weak) control Lyapunov function-based nonlinear receding horizon control, with randomized optimization. This approach is applied to the problem of robot navigation in the presence of state and input constraints. It is shown that under certain conditions, relaxing the definiteness requirements on the terminal cost function allows one to select control inputs through a Monte-Carlo optimization scheme in a way that preserves the stability and convergence properties of the closed loop system. While the particular randomized optimization scheme used here can be substituted for the nonlinear optimal control method of choice, the introduction of randomization in receding horizon optimization is anticipated to offer additional trade-offs between performance and computation speed compared to the fixed-overhead nonlinear optimal control strategies typically employed.

Index Terms—Model predictive control; randomized algorithms; robot navigation

I. INTRODUCTION

In the motivating application, a robot moves between any two configurations in its workspace, under constraints on its state and inputs, while attempting to optimize some mission performance metrics. The control strategy can be switching, with the robot dynamics given as

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

where $x \in \mathbb{R}^n$ is the state, $u \in \mathbb{R}^m$ the control input, and $\sigma \in \Sigma \subseteq \mathbb{R}$ is a *controlled* switching signal indicating the particular dynamics currently active. The robot's workspace is a set $\mathcal{W} \subset \mathbb{R}^n$ and the system may be constrained to stay clear of the union of possibly disconnected regions of the state space $\mathcal{O} \subset \mathcal{W}$ representing obstacles or state constraint sets. The free workspace is then the set $P \triangleq \mathcal{W} \setminus \mathcal{O}$, and the boundary of P is assumed known.

The approach outlined in this note is a blend of control Lyapunov function-based model predictive control schemes, with navigation functions and randomized optimization algorithms. What motivates the search for a solution with these ingredients is the desire to deploy miniature mobile robots which execute optimal or near-optimal motion plans with guaranteed stability and convergence properties. 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 [2]. The motion planning nature of the task, as well as the need for a control Lyapunov function, suggests the integration of navigation functions into the nonlinear receding horizon framework. Randomized algorithms can then be used to simplify the online optimization task, and enforce task-specific constraints.

The use of control Lyapunov functions as terminal costs in receding horizon nonlinear optimization offers closed loop stability in the infinite horizon [2]–[4], as well as robustness [5]. Enforcing stability by means of terminal constraints [6] imposes additional computational overhead. For the unconstrained case, an appropriately long control

Herbert Tanner is with the Department of Mechanical Engineering at the University of Delaware, Newark DE 19716. Email: btanner@udel.edu. Fax: (302) 831 3619. Jorge Piovesan is with K&A Wireless LLC, 2617 Juan Tabo NE Suite A, Albuquerque NM 87112. Email: jpiovesan@ka-wireless.com. A preliminary version of this paper has appeared in [1].

This work is supported by ARL through a subcontract from the University of New Mexico as part of MAST CTA #W911NF-08-2-0004.

horizon may permit excluding a terminal cost in the objective function [5], although still useful for numerical conditioning [2].

Randomized algorithms are an alternative to exact analytic or numerical optimization techniques [7], [8]. The (NP-hard) complexity of several robust stability design problems [9], motivates the use of randomized algorithms for control synthesis [10]–[12]. These randomized approaches offer explicit bounds on the number of samples required to guarantee specific confidence and accuracy levels.

Robot motion planning in the presence of obstacles is also an NP-hard problem [13]. To address such problems both purely analytical [14], numerical [15], randomized [16], as well as hybrid [17], [18] methods have been proposed, with the latter aimed at addressing the problem of local minima in the classical potential field approach. In one of the several available approaches to deal with the problem of local minima [19], navigation functions are introduced as a special class of potential functions which can be tuned¹ so that the only minimum is the desired configuration. An additional 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 impose limitations on the use of navigation functions, since the gradient of latter cannot be guaranteed to be non-zero, except for trivial cases.

The technical challenge addressed in this note 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. Randomization compensates for inadequate tuning of navigation function parameters, and allows performance/speed trade-offs in receding horizon by regulating computational overhead in each loop. Receding horizon control, on the other hand, preserves the convergence of solutions in navigation function-based motion planning and allows the simultaneous satisfaction of secondary objectives and constraints.

II. NAVIGATION FUNCTIONS AND SWITCHING DYNAMICS

Consider a mechanical system with $p \in P$ being the vector of its generalized coordinates. A navigation function can then be defined as a total map $\varphi : P \rightarrow [0, 1]$, [19] which is smooth (or at least C^2), has a unique minimum at a single point $p_d \in P$, is uniformly maximal on the boundary of P , and is a Morse function. It has been shown [20] that the following construction

$$\varphi(p) = \frac{\|p - p_d\|^2}{(\|p - p_d\|^{2\kappa} + \beta)^{1/\kappa}}, \quad (2)$$

where β is the product of functions $\beta_i = \|p - p_i\| - r_i$, each representing a spherical obstacle of radius r_i centered at p_i in the robot's workspace, can be turned into a navigation function for a sufficiently large value of $\kappa > 0$. For cases where the system dynamics can be adequately described by $\dot{p} = u$, setting $u = -\nabla_p \varphi(p)$ stabilizes the system at p_d while avoiding the regions where β_i are negative. In that sense, the navigation function is practically a (local) control Lyapunov function [21] for the system with state $x = (p, \dot{p})^T$, that is, is a differentiable and positive definite function $V : P \rightarrow [0, 1]$ for which $\inf_u \left\{ \frac{\partial V(x)}{\partial x} f(x, u) \right\} < 0, \forall x \neq 0$. For $\varphi(p)$, this condition is satisfied almost everywhere in P , with the exception of a set of measure zero consisting of the saddle points of φ . These points might appear as a small nuisance in the case of continuous-state feedback,

¹Although 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.

but they present a challenge in the discontinuous case, when the condition on the Lyapunov function is strengthened [22]–[24] in the form $\inf_u \left\{ \frac{\partial V(x)}{\partial x} f_\sigma(x, u) \right\} \leq -W(x)$, with $W(x)$ being a continuous positive definite function. This condition cannot be satisfied when navigation functions are used as control Lyapunov functions, unless one is prepared to exclude neighborhoods of the saddle points. This motivates one to seek alternative conditions for asymptotic stability for the case of switching dynamics or discontinuous feedback, which do not bound the derivative of the control Lyapunov function below another negative definite function.

Suppose that the switching in (1) is caused exclusively due to a time-invariant feedback control law $u = k(\sigma; x)$. The closed loop dynamics are then written as $\dot{x} = f_\sigma(x)$, and asymptotic stability can be guaranteed under the conditions of the following result:

Proposition 1: Let $V(x)$ be a differentiable, positive definite function that is lower bounded by a class- \mathcal{K}_∞ function $\theta(\|x\|)$. If for a class- \mathcal{K}_∞ function $\gamma(\cdot)$, within each of the switching intervals $(\tau_i, \tau_{i+1}]$ in a dwell-time switching signal σ ,

$$\max_{t \in (\tau_i, \tau_{i+1}]} \left\{ \frac{\partial V}{\partial x} f_\sigma(x) \right\} < \frac{\theta(\|x(\tau_i)\|)}{\tau_{i+1} - \tau_i}, \quad (3)$$

$$\int_{\tau_i}^{\tau_{i+1}} \frac{\partial V}{\partial x} f_\sigma(x) dt \leq -\gamma(\|x(\tau_i)\|), \quad (4)$$

then $\dot{x} = f_\sigma(x)$ is asymptotically stable at the origin.

Proof: Pick an $\varepsilon > 0$ and define a ball $B_\varepsilon \triangleq \{x \in \mathbb{R}^n \mid \|x\| \leq \varepsilon\}$. 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 $\{x \in \mathbb{R}^n \mid V(x) \leq V_\varepsilon\}$, 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 $\{x \in \mathbb{R}^n \mid V(x) \leq V_\delta\}$ containing the origin. Fit a ball B_δ inside Ω_δ , by defining it as $B_\delta = \{x \in \mathbb{R}^n \mid \|x\| \leq \min_{x \in \partial \Omega_\delta} \|x\|\}$. Assume that at time τ_i , $x(\tau_i) \in B_\delta$. Since $B_\delta \subseteq \Omega_\delta$, $V(x(\tau_i)) \leq V_\delta = V_\varepsilon - (\tau_{i+1} - \tau_i)d$, where $d \triangleq \frac{\theta(\|x(\tau_i)\|)}{\tau_{i+1} - \tau_i}$. 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 (3), $V(x(t)) \leq V(x(\tau_i)) + d(t - \tau_i)$. Thus, $x(t) \in \Omega_\varepsilon$, for all $t \in (\tau_i, \tau_{i+1}]$. In addition, due to (4) $x(\tau_{i+1}) \in \Omega_\delta$, and an inductive argument using (3) and (4) shows that $x(\tau_k) \in \Omega_\delta \subset \Omega_\varepsilon$ for all $k \geq i$. 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 (4) implies the existence of a converging sequence $\{V(x(\tau_k))\}$, since $V(x)$ is lower bounded by 0 away from $x = 0$ and it is strictly decreasing. It follows that as $k \rightarrow \infty$, the difference between two consecutive terms $\|V(x(\tau_i)) - V(x(\tau_{i+1}))\|$ should converge to zero. Unless $\lim_{k \rightarrow \infty} V(x(\tau_k)) = 0$, this means that $\gamma(\|x(\tau_k)\|) \rightarrow 0$, without $\|x(\tau_k)\| \rightarrow 0$, which is impossible since $\gamma(\cdot)$ is a class- \mathcal{K}_∞ function. Therefore, $\{V(x(\tau_k))\} \rightarrow 0$ as $k \rightarrow \infty$ implying $\|x(\tau_k)\| \rightarrow 0$ when $k \rightarrow \infty$. Thus, the origin is asymptotically stable. ■

Compared to [25], Proposition 1 effectively bounds the worst case increase of V within the switching interval by a class- \mathcal{K}_∞ function of state and dwell time, which forces a predictable and faster rate of convergence at the origin. It should be noted that in the context of Proposition 1, the switching signal σ is thought to be controllable: the designer may regulate the length of the switching intervals to satisfy (4)-(3).

III. MODEL PREDICTIVE NAVIGATION

In finite horizon optimization, one aims to minimize the functional

$$J_T(x, u_T(\cdot)) \triangleq \int_0^T q(x^u(\tau; x), u_T(\tau; x)) d\tau + V(x^u(T; x)), \quad (5)$$

that quantifies the cost of flowing along a closed loop system trajectory $x^u(t; x)$ starting at x , under the control law $u_T(t; x)$, for a time interval in which $t \in [0, T]$. The function $V(\cdot)$ in (5) is an approximation of the tail of the infinite horizon integral, now truncated at T . The problem is to determine the optimal control law u_T^* , which minimizes the finite horizon cost from x , $u_T^*(x; \cdot) \triangleq \arg \min J_T(x, u_T(\cdot))$, giving rise to an optimal finite horizon trajectory $(x_T^*(t; x), u_T^*(t; x))$, for $t \in [0, T]$, with x_T^* denoting the optimal closed loop trajectory for that time interval. In a receding horizon strategy, one uses $u_T^*(t; x)$ for $t \in [0, \zeta]$, with $\zeta < T$, and then recomputes for a new initial state $x^u(t; x_T^*(\zeta, x))$.

Control Lyapunov functions can be used to approximate the terminal cost within a finite horizon optimization [3], [4]. Although relaxed stability conditions have been suggested [5], the standard assumption for the control Lyapunov function in this framework [2] is that it is compatible with the incremental cost $q(x, u)$, in the sense that $\inf_{u(\cdot)} \left(\frac{\partial V}{\partial x} f(x, u) + q(x, u) \right) \leq 0$. A navigation function cannot satisfy this compatibility in any small neighborhood of its critical (saddle) points.

The assumption that follows is used as a fact in [2], because it can be guaranteed based on the continuity and properness of the infinite and finite horizon cost functions. We choose to state it as a working assumption to place emphasis on it, because it is pivotal in establishing asymptotic stability. Here, Ω is defined 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_∞ for which (i) $J_\infty^*(x) \geq m_\infty \|x\|^2$, and (ii) $J_T^*(x) \leq c_T \|x\|^2$.

The next assumption is in fact a stability condition, and it is the equivalent of the compatibility condition [2, Eqn. (2.4)] for the case considered here, where a control Lyapunov function is not available:

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

$$\int_0^\zeta q(x(\tau), k(x(\tau))) d\tau \leq \gamma(\|x(0)\|) \leq - \int_0^\zeta \dot{V}(x(\tau), k(x(\tau))) d\tau, \quad \forall x(0) \in \Omega. \quad (6)$$

Based on these two assumptions, we can state the following results.

Proposition 2: [cf. [2, Theorem 2.2.1]] Suppose that $x \in \mathbb{R}^n$ and $T > 0$ are such that $x_T^*(T; x) \in \Omega$. Then, for any $\zeta(x)$, and feedback control $u = k(x(t))$ for which there exists a \mathcal{K}_∞ class function γ satisfying (6) for $x \in \Omega$, the optimal cost from $x_T^*(\zeta; x)$ is such that

$$J_T^*(x_T^*(\zeta; x)) \leq J_T^*(x) - \int_0^\zeta q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau. \quad (7)$$

Proof (sketch): 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 $x_T^*(\zeta; x)$, and expand it so that $\int_T^{T+\zeta} q(\tilde{x}(\tau), \tilde{u}(\tau)) d\tau + \int_T^{T+\zeta} \dot{V}(\tilde{x}(\tau)) d\tau$ appear explicitly. Noting that the aforementioned term is nonpositive due to (6), $J_T(x_T^*(\zeta; x), \tilde{u}(\cdot))$ is upper bounded by $J_T^*(x) - \int_0^\zeta q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau$. The result follows since the optimal cost satisfies $J_T^*(x_T^*(\zeta; x)) \leq J_T(x_T^*(\zeta; x), \tilde{u}(\cdot))$. ■

Although (3) is not explicitly invoked in the following discussion, it has to be enforced to ensure stability, and satisfaction of the constraints encoded in $V(x)$ when the latter is constructed as a navigation function. Specifically, $\theta(\|x_0\|)$ must be small enough so

that $V(x(t)) + \theta(\|x_0\|) < 1$, for all $t \in [0, \zeta]$. Also note that ζ 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 (6) for a sequence of control horizons $\{\zeta_\ell\}$, with terms defined recursively as $\zeta_\ell = \zeta(x(\sum_{k=1}^{\ell-1} \zeta_k))$. Then (7) still holds.

Proof (sketch): The argument follows almost inductively, based on the proof of Proposition 2. In this case it is shown that $J_T(x_T^*(\zeta_1; x), \tilde{u}(\cdot)) = J_T^*(x) - \int_0^{\zeta_1} q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau$, and the proof is completed by noting that $J_T^*(x_T^*(\zeta; x)) \leq J_T(x_T^*(\zeta; x), \tilde{u}(\cdot))$. ■

Lemma 1: [cf. [2, Lemma 2.2.1]] Suppose that $x \in \Omega$. Then $x_T^*(T; x) \in \Omega$ for any $T = \sum_{i=1}^{\ell} \zeta_i$, and for every $\ell \geq 1$, where ζ_i satisfies (6).

Proof: The statement follows directly from the stability argument in the proof of Proposition 1, once inequality (6) is established. ■

Proposition 3: [cf. [2, Prop 2.2.2]] 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 $(x_{rh}(t), u_{rh}(t))$, $t \geq 0$, resulting from the use of a receding horizon strategy $\mathcal{RH}(T, \{\zeta_k\})$ with $\delta_k > 0$, $\sum_{k=0}^{\ell} \zeta_k \rightarrow \infty$ as $\ell \rightarrow \infty$. Then, $J_\infty(x_0, u_{rh}(\cdot)) \leq J_T^*(x_0)$.

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 $J_T^*(x_k) - J_{T-\zeta_{k+1}}^*(x_{k+1}) \geq 0$. The total cost of this strategy satisfies $J_\infty(x_0, u_{rh}(\cdot)) \leq J_T^*(x_0)$, because $J_T^*(x_k) \leq J_{T-\zeta}^*(x_k)$, for all $\zeta \geq 0$ and $k \geq 0$. ■

Theorem 1: [cf. [2, Thm 2.2.2]] Let $T > 0$ and consider the use of a receding horizon scheme $\mathcal{RH}(T, \{\zeta_k\})$ with each $\zeta_k \in (0, T]$ and $\sum_{j=0}^k \zeta_j \rightarrow \infty$ as $k \rightarrow \infty$. Then, for each $x_0 \in \Gamma$, the resulting trajectory converges to the origin exponentially fast.

Proof: Define the continuous function $W(t; x_0, u_{rh}(\cdot)) \triangleq \int_t^\infty q(x_{rh}(\tau), u_{rh}(\tau)) d\tau$. By Proposition 3, $W(t; x_0, u_{rh}(\cdot)) \leq J_T^*(x_{rh}(t))$, and since $\mathcal{RH}(T, \{\delta_k\})$ is suboptimal, $J_\infty^*(x_{rh}(t)) \leq W(t; x_0, u_{rh}(\cdot))$. Now, $\frac{\partial}{\partial t} W(t; x_0, u_{rh}(\cdot)) \leq -\frac{c_T}{c_T} W(t; x_0, u_{rh}(\cdot))$, and using the bounds of Assumption 1 from the Comparison Lemma it follows $m_\infty \|x_{rh}(t)\|^2 \leq c_T e^{-\frac{c_T}{c_T} t} \|x_0\|^2$, which implies that $\|x_{rh}(t)\|$ converges exponentially to the zero. ■

Note 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 (3) and (4).

IV. RANDOMIZED MODEL PREDICTIVE CONTROL

This section suggests an alternative approach to finite-time nonlinear optimal control design using randomized algorithms. Instead of numerically computing the finite horizon optimal control law $u_T^*(\cdot)$ one can evaluate *probable near minimizers* of the finite-horizon cost J_T , using Monte Carlo sampling. Simple probabilistic reasoning offers (conservative) lower bounds on the sample size which relate to the volume of the solution space which is missed by the sampling.³

Given that a nominal, feasible and stabilizing solution is available (the one for which the conditions of Proposition 1 are satisfied), the

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

³The sampling algorithm used here is a placeholder for the randomized algorithm of choice. Its use is motivated by our interest to extend the proposed method to the case of uncertainty.

randomized algorithm returns an alternative solution that improves on the performance of the nominal one.

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 controls satisfying (6) and $J_T(x_{i+1}, u_{i+1}(\cdot)) < J_{T-\zeta_i}(x_{i+1}, u_i(\cdot + \zeta_i))$, for all $i \in \mathbb{N}_+$. Then $x^{u_i}(t) \rightarrow 0$ as $t \rightarrow \infty$.

Proof: Define the sequence of finite horizon costs $\{c_i\}_{i=1}^{\infty}$, with $c_i \triangleq J_T(x_i, u_i(\cdot))$, and observe that for $T = \zeta_i$, $c_i - c_{i+1} > \int_0^{\zeta_i} q(x^{u_i}(\tau; x_i), u_i(\cdot)) d\tau$. Condition (6) 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, then $c_i - c_{i+1} \rightarrow 0$, while $\gamma(\|x_i\|) \rightarrow c > 0$, which is a contradiction. Therefore one has to have $c_i \rightarrow 0 \Rightarrow q(x, u) \rightarrow 0 \Rightarrow x \rightarrow 0$. ■

Let each feedback control law $u = k(x)$ be parameterized in terms of a set of control parameters $\eta \in H$. Thus, given a state x , the finite horizon cost becomes a function of the parameter (decision) vector η . If the parameterization chosen is complete, 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 $J_T(x, \eta) = J_T^*(x)$.

Assume a probability measure \mathbb{P} on H . It is straightforward to show that if μ_i , $i = 1, \dots, N_s \geq \frac{\ln(1/\delta)}{\ln(1/(1-\alpha))}$ independent and identically distributed random samples are drawn from H according to a probability distribution $\mathcal{P}(\eta)$, where α and δ are positive parameters in the interval $(0, 1)$, then $J_T^\circ \triangleq \min_i J(\eta_i)$ is a *probable near minimum* [8], [11], [26] of J_T^* to level α and confidence $1 - \delta$, in the sense that for every set $\tilde{H} \subset H$ measuring $\mathbb{P}(\tilde{H}) \leq \alpha$, one has $\mathbb{P}\left\{\inf_{H} J_T(\eta) \leq J_T^\circ \leq \inf_{\tilde{H}} J_T(\eta)\right\} \geq 1 - \delta$.

Algorithm 1 Randomized optimization

Require: $\alpha, \delta, J_T, P, H, T, \zeta, \gamma, \theta, x^\eta, V, \dot{V}$.

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 $\mathcal{P}(\eta)$

5: $\tilde{\eta}^+ \leftarrow \{\eta_i \mid V(x^\eta(\zeta, x)) - V(x) \leq \gamma(\|x\|) \wedge \max_{t \in [0, \zeta]} |V(\eta_i; x^\eta(t, x))| \leq \frac{\theta(\|x\|)}{\zeta}\}$

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: $\eta^\circ \leftarrow \arg \min_{\tilde{\eta}} J_T(\eta_i)$


The algorithm will halt at step 3 if the stability constraints of Proposition 1 are not satisfied for the number of samples determined by the choice of α and δ . Variations of the above algorithm could include an adaptive increase in the number of samples until an upper bound is reached, which can be associated with the largest computational cost the control designer is willing to accept. There will always be a trade-off between performance and the ability to satisfy the stability conditions on one hand, and the corresponding computational cost on the other.

V. NUMERICAL RESULTS


The simulation described here is a navigation scenario for an agent with single integrator dynamics $\dot{p} = u$ moving in a two dimensional sphere-world environment (cf. [19]). Here, $x = (p, \dot{p})^T$, and the environment topology captured by a function φ of the form (2), which is not tuned properly to become a navigation function (Fig. 1). As a result, a local minimum with nontrivial attraction basin exists.

In addition to moving from initial to final configuration, the agent also needs to minimize the shadowing effect of obstacles to the

radio signal from a stationary base station, captured by a positive semi-definite cost function $C(x)$. The combined cost function for the receding horizon optimization takes the form $J_T(x, u) = \int_0^T \{c_1 \|u(\tau)\|^2 + c_2 \|x(\tau)\|^2 + c_3 C(x(\tau))\} d\tau + \varphi(x)$, where c_i , $i = 1, \dots, 3$ are constants. The input u is parameterized in terms of a polynomial function basis, and sampling is performed on the space of the constant coefficients.



(a) The initial position for the agent is marked by a hollow dot. Its location after ten cycles is shown by a filled dot and its path is traced by a solid curve. The two circular obstacles are located at coordinates $(-0.2, \pm 0.3)$, while the radio base station is at $(0, 0.3)$. The agent starts from within the attraction region of the local minimum (within the closed contour around $(0.7, 0)$) and needs to stabilize at the origin.



(b) The evolution of optimization costs over time. The change in the rate of increase of the incremental cost at around 0.6 sec is due to the agent establishing line-of-sight with the base station. Note also the slight, short-term increase in the terminal cost around the same time—possibly due to the maneuver for obtaining better signal. This temporary increase is permitted by Proposition 1.

Fig. 1. The communication-aware navigation task for a mobile agent is performed using randomized receding horizon optimization.

Typical steepest descent along the gradient of ϕ results in the agent being trapped at the local minimum configuration. In our numerical tests, however, there were consistently samples that steered the agent away from the basin, although the randomized nature of the design offers no guarantee that such samples will *always* be found. Randomization cannot always compensate for a bad design.

VI. CONCLUSIONS

Lyapunov-based conditions for asymptotic stability of switching nonlinear systems can be relaxed, so that stability and convergence of solutions in receding horizon optimization can be guaranteed with a terminal cost that does not have a negative semi-definite derivative. Instead, the CLF-like condition can be replaced by an integral inequality ensuring that the function decreases significantly between switching times. This fact alleviates the consequences on convergence of systems steered using inadequately tuned navigation functions, and overcomes a limitation to their use as receding horizon terminal cost functions. The application of a randomized algorithm suggests an alternative to the use of Newton-based nonlinear optimization algorithms in model predictive control while preserving the stability and convergence properties of the receding horizon strategy. It is conceivable that in a future implementation randomized and gradient-based trajectory optimization are combined; one possible way would be to locally optimize each sample input trajectories using projection operators [27], an approach which will also ensure that the inputs tested for optimality are all feasible.

REFERENCES

- [1] J. L. Piovesan and H. G. Tanner, "Randomized model predictive control for robot navigation," in *IEEE International Conference on Robotics and Automation*, 2009, pp. 94–99.
- [2] A. Jadbabaie, "Receding horizon control of nonlinear systems: A control lyapunov function approach," Ph.D. dissertation, California Institute of Technology, Pasadena, California, October 2000.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] S. Keerthi and E. Gilbert, "Optimal infinite-horizon feedback laws for a general class of constrained discrete time systems: Stability and moving-horizon approximations," *Journal of Optimization Theory and Applications*, vol. 57, no. 2, pp. 265–293, 1988.
- [7] M. Vidyasagar, *A Theory of Learning and Generalization*, ser. Communications and Control Engineering. London, UK: Springer, 1997.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] M. Vidyasagar, "Randomized algorithms for robust controller synthesis using statistical learning theory," *Automatica, Elsevier Science Ltd.*, vol. 37, pp. 1515–1528, 2001.
- [12] 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.
- [13] J. F. Canny, *The Complexity of Robot Motion Planning*. MIT Press, 1998.
- [14] 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.
- [15] 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.
- [16] L. Kavrakı, 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.

- [17] 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.
- [18] H. Hagenaaars, J. Imura, and H. Nijmeijer, "Approximate continuous-time optimal control in obstacle avoidance by time/space discretization of non-convex state constraints," in *IEEE International Conference on Control Applications*, 2004, pp. 878–883.
- [19] 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.
- [20] D. E. Koditschek and E. Rimon, "Robot navigation functions on manifolds with boundary," *Advances in Applied Mathematics*, vol. 11, pp. 412–442, 1990.
- [21] 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.
- [22] D. Liberzon, *Switching in Systems and control*. Boston, MA, USA: Birkhauser, 2003.
- [23] 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.
- [24] 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.
- [25] H. Ye, A. N. Michel, and L. Hou, "Stability theory for hybrid dynamical systems," *IEEE Transactions on Automatic Control*, vol. 43, no. 4, pp. 461–474, 1998.
- [26] R. Tempo, E. Bai, and F. Dabbene, "Probabilistic robustness analysis: Explicit bounds for the minimum number of samples," *Systems and Control Letters*, vol. 20, pp. 237–242, 1997.
- [27] J. Hauser, "A projection operator approach to optimization of trajectory functionals," in *IFAC World Congress*, Barcelona, Spain., 2002.