POLITECNICO DI TORINO
Repository ISTITUZIONALE

A Finite-Gain Stability Approach to NMPC Design

Original
A Finite-Gain Stability Approach to NMPC Design / Novara, C., Boggio, M., Calogero, L., Pagone, M.. - ELETTRONICO. - (2026). (65th Conference on Decision and Control (CDC) Honolulu (USA) 15-18 Dicembre, 2026).

Availability:
This version is available at: 11583/3015670 since: 2026-09-17T10:48:42Z

Publisher:
IEEE

Published
DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright
IEEE postprint/Author's Accepted Manuscript

©2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

26 November 2026

A Finite-Gain Stability Approach to NMPC Design

Carlo Novara, Mattia Boggio, Lorenzo Calogero, Michele Pagone

***Abstract*— This paper proposes a novel approach to the design of Nonlinear Model Predictive Control (NMPC) schemes based on the Finite-Gain Stability (FGS) concept. The proposed formulation considers the case where the plant is affected by unknown but bounded disturbances, which renders the classical Lyapunov-based analysis/design difficult. Based on FGS conditions for a closed-loop system, we develop a systematic *scenario-based* NMPC design methodology, allowing us to choose the relevant NMPC parameters enforcing FGS and providing satisfactory tracking performance. A simulated example is presented to demonstrate the effectiveness of our framework.**

I. Introduction

In the framework of optimal control, Nonlinear Model Predictive Control (NMPC) has emerged as one of the most effective and reliable control methodologies for nonlinear systems. Its success stems from its inherent ability to provide optimal control commands for multi-variable systems in the presence of inputs, outputs, and state constraints [1].

In this context, establishing conditions for closed-loop stability is a nontrivial problem that poses exceptional challenges. This problem is even more relevant in the presence of disturbances, whose effects on the system state and tracking error are difficult to characterize. It is then natural to ask under which conditions the *nominal* NMPC[1] is able to ensure that perturbed trajectories remain bounded or converge close to the desired reference. Such a problem is, in general, difficult to characterize.

Nowadays, nominal NMPC stabilizing design theory is typically based on the Lyapunov approach, see, e.g. [2], and the references therein. Nevertheless, its extension to the case where the plant is affected by unknown disturbances and time-varying references is nontrivial. Early works (see, e.g., [3], [4]) provide Lyapunov-based asymptotic stability certificates without, however, offering constructive guidelines for determining disturbance bounds and for explicitly quantifying their effect on the tracking performance, making the stabilizing design rather conservative and hard to apply in practice.

Alternative approaches, which can be more convenient than the standard Lyapunov-based procedure in certain settings, formulate the closed-loop stability of nonlinear systems by leveraging the notions of Input-to-State Stability (ISS) [5] and Finite-Gain Stability (FGS) [6], [7]. Both these concepts characterize how the system state/output converges or remains bounded, depending on the initial condition, the magnitude of the reference, and the bounds on the disturbance.

Building up on the recent results on ISS and FGS, in this paper, we propose a novel FGS-based design approach to NMPC control systems. We establish sufficient conditions under which, given a set of feasible initial conditions and possibly time-varying references, under the application of the NMPC control law, the closed-loop trajectories remain bounded and converge towards a neighborhood centered on the reference.

The proposed framework is not meant to serve as a robust counterpart of the nominal NMPC (as in [5], [8]), rather, it defines conditions under which a NMPC controller is inherently capable to attenuate the external disturbance for letting the system's output to converge and remain bounded in a close neighborhood of the reference. The stability of the closed-loop is not formulated by employing the standard terminal ingredients of the Lyapunov framework [2], but relies on the FGS bound estimates. However, unlike the classical ISS-based and FGS-based methods, see, e.g., [9], thanks to a quasi-LPV (Linear Parameter Varying) formulation, we explicitly deal with the case of time-varying references, but in a different fashion with respect to the artificial reference tracking MPC approach, see, e.g., [10], [11] and the references therein. Indeed, in this latter approach, the terminal constraints/penalties may be hard to design for arbitrary time-varying references. Instead, we formulate FGS concepts using a tracking-oriented error bound that directly accounts for reference variations.

Afterwards, we focus on the choice of the NMPC configuration parameters, namely the prediction horizon and the weighting matrices appearing in the cost function. Given a region of initial conditions, a class of reference signals, the goal is to provide a scenario-based selection of NMPC configurations that enforces FGS when the plant is subject to unknown but bounded disturbances.

To this end, we propose an FGS-based three-step offline NMPC design procedure: i) a finite simulation horizon is selected so as to capture the transient interval over which the contraction of the closed-loop error dynamics becomes visible; ii) a family of candidate NMPC configurations is generated by varying the main tuning parameters around a nominal design; iii) each candidate is evaluated over a

The work of C. Novara and M. Pagone was supported by the Space It Up project, funded by the Italian Space Agency (ASI) and the Ministry of University and Research (MUR) under contract n. 2024-5-E.0, CUP n. I53D24000060005. The authors are with the Department of Electronics and Telecommunications, Politecnico di Torino, 10129 Turin, Italy ({carlo.novara, mattia.boggio, lorenzo.calogero, michele.pagone}@polito.it)

[1]Henceforth, with *nominal* NMPC we shall denote a *non-robust* controller which, in principle, is still able to guarantee the stability of the closed-loop in presence of external disturbance.

scenario set composed of initial conditions and reference signals. In this way, the FGS analysis is used as a practical guide for controller synthesis.

The appeal of the proposed methodology stems from the fact that each candidate NMPC configuration is assessed only over a finite horizon, using quantities that are meaningful both from a stability and from a performance viewpoint. This makes the procedure computationally efficient and particularly suitable for the practical tuning of NMPC controllers over a prescribed operating domain, while preserving a high level of interpretability.

The novel contribution of the paper is twofold: i) we derive a formulation of the finite-gain bounds for the closed-loop tracking error; ii) we exploit such bounds to formulate a scenario-based offline design methodology for selecting an NMPC configuration enforcing the closed-loop FGS.

The proposed approach is illustrated on an automated-vehicle control problem involving lateral and longitudinal motion, where the method is used to choose the prediction horizon and cost weights before validating the resulting controller on lane-keeping and obstacle-avoidance maneuvers.

II. Discrete-time NMPC

Consider a discrete-time nonlinear system described by the following state equation:

$$x_{t+1} = f(x_t, u_t) + d_t \tag{1}$$

where $t \in \mathbb{N}_0$ is the time index, $x_t \in \mathbb{R}^{n_x}$ is the system state, $u_t \in U \subset \mathbb{R}^{n_u}$ is the command input, $d_t \in D \subset \mathbb{R}^{n_d}$ is an exogenous disturbance, U and D are compact sets. The state is assumed to be measured in real-time. If not measured, it can be estimated by means of an observer or an input-output model can be used instead of the state equations. The control task is to make the system state x_t track a desired reference signal $r_t \in \mathbb{R}^{n_x}$.

The NMPC formulation developed in this paper is as follows. At each time step t, the system state x_t is measured. An optimal command is computed based on x_t, and then applied to the system. Let $T \geq 1$ be the prediction horizon, the NMPC objective function is defined as

$$J_o(x_t, \mathbf{u}) \doteq \sum_{k=0}^{T-1} \mathfrak{u}_k^\top R \mathfrak{u}_k + \sum_{k=1}^{T-1} \tilde{x}_k^\top Q \tilde{x}_k + \tilde{x}_T^\top P \tilde{x}_T \tag{2}$$

where $\tilde{x}_k \doteq \mathfrak{r}_k - \hat{x}_k$ is the predicted tracking error, $\mathfrak{r}_k \doteq r_{t+k} \in \mathbb{R}^{n_x}$, $k = 1, \dots, T$, is the reference to track, $\mathbf{u} \doteq (\mathfrak{u}_0, \dots, \mathfrak{u}_{T-1}) \in \mathbb{R}^{Tn_u \times 1}$ is the applied input sequence, and R, Q, P are diagonal positive-definite weight matrices.

The input constraints are described by $\mathfrak{u}_k \in U$, $k = 1, \dots, T$. The state constraints are treated as *soft constraints* using penalty functions. In particular, for each constraint, a penalty function $\sigma_j \in \mathcal{C}^1$ is introduced, such that $\sigma_j(\hat{x}_k) \simeq 0$ when the constraint is satisfied, $\sigma_j(\hat{x}_k) \gg J_o(x_t, \mathbf{u})$ when it is not (see, e.g., [12]). The penalty functions are added to J_o, giving the "augmented" objective function

$$J(x_t, \mathbf{u}) \doteq J_o(x_t, \mathbf{u}) + \textstyle\sum_{j,k} \sigma_j(\hat{x}_k) \ . \tag{3}$$

Moreover, the command sequence $\mathbf{u} \doteq (\mathfrak{u}_0, \dots, \mathfrak{u}_{T-1})$ may contain a relatively large number of decision variables. An efficient method to reduce the number of variables is the Move Blocking technique, see, e.g., [13], [14]: the command sequence $\mathbf{u}$ is parametrized as $\mathbf{u} = \Gamma$, where $\Gamma \in \mathbb{R}^{Tn_u \times n_c}$ is a full-rank matrix with $n_c < Tn_u$, and $c \in \mathbb{R}^{n_c}$ is a new command input sequence with reduced dimension. A sample $\mathfrak{u}_k$ of $\mathbf{u}$ can be obtained using a selection matrix S_k^u such that $\mathfrak{u}_k = S_k^u \mathbf{u} = S_k^u \Gamma c$. Note that Γ is a design parameter, selected a-priori. Under the above input parametrization, the optimization problem becomes:

$$\begin{aligned} c^* = \arg\min_c \ & J(x_t, \Gamma c) \\ \text{subject to:} \\ & \hat{x}_0 = x_t \\ & \hat{x}_{k+1} = f(\hat{x}_k, S_k^u \Gamma c), \ k = 0, \dots, T-1 \\ & S_k^u \Gamma c \in U, \ k = 0, \dots, T-1. \end{aligned} \tag{4}$$

As usual in the NMPC context, the plant is controlled by NMPC under the one-step receding horizon policy.

III. Finite-gain analysis

In this section, a closed-loop finite-gain stability analysis is developed for the case where an exact model of the plant (1) is available.

The finite-gain stability concept considered is the following.

Definition 1. *Consider a discrete-time nonlinear system $x_{t+1} = g(x_t, w_t)$, where $x_t \in \mathbb{R}^{n_x}$ is the state, $x_0 \in X_0 \subseteq \mathbb{R}^{n_x}$ is the initial condition, and $w_t \in W \subseteq \mathbb{R}^{n_w}$ is the input (containing all external signals acting on the system, including commands and disturbances). The system is finite-gain stable on $(X_0, \mathcal{W})$ if a constant $0 \leq \rho < 1$, and finite non-negative constants $\Lambda_0, \Lambda_w, \Lambda_b$ exist, such that*

$$\|x_t\| \leq \Lambda_0 \rho^t \|x_0\| + \Lambda_w \sum_{i=0}^{t-1} \rho^{t-i} \|w_i\| + \Lambda_b, \quad t = 1, 2, \dots$$

for all initial conditions $x_0 \in X_0$ and input sequences $(w_0, w_1, \dots) \in \mathcal{W}$. Stability is said without bias if $\Lambda_b = 0$. □

A finite-gain stability analysis is now developed for the closed-loop system given by the plant (1) controlled in feedback by the NMPC strategy of Section II. This system is described by $x_{t+1} = f(x_t, u_t) + d_t$, where $u_t = S_1^u \Gamma c^*$ and c^* is the solution of problem (4) at time t. Since this solution depends on x_t and the reference sequence $\mathbf{r}_t \doteq (r_{t+1}, \dots, r_{t+T})$, we use the notation $c^* \equiv c^*(x_t, \mathbf{r}_t)$. The closed-loop system state equation is

$$\begin{aligned} x_{t+1} &= f(x_t, u_t) + d_t = f(x_t, S_1^u \Gamma c^*(x_t, \mathbf{r}_t)) + d_t \\ &= f_c(x_t, \mathbf{r}_t) + d_t \end{aligned} \tag{5}$$

where $f_c(x_t, \mathbf{r}_t) \doteq f(x_t, S_1^u \Gamma c^*(x_t, \mathbf{r}_t))$ is the closed-loop transition function. The initial state and each reference sample belong to a set where the penalty functions are not active:

$x_0 \in X_0 \doteq \bar{X}_0 \cap X_{\rm na}$, where $\bar{X}_0 \subset \mathbb{R}^{n_x}$ is a compact connected set and $X_{\rm na} \doteq \{z \in \mathbb{R}^{n_x} : \sigma_j(z) = 0, \forall j\}$. Similarly: $r_t \in R \doteq \bar{R} \cap X_{\rm na}$, where $\bar{R} \subset \mathbb{R}^{n_x}$ is a compact connected set. The reference sequence $\boldsymbol{r} = (r_0, r_1, \dots)$ belongs to a signal set of interest: $\boldsymbol{r} \in \mathcal{R} \subseteq R^\infty$, where $\mathcal{R} \doteq \{(r_0, r_1, \dots) : r_k \in R\}$ and exogenous disturbance sequence $\boldsymbol{d} = (d_0, d_1, \dots)$ belongs to the bounded set $\mathcal{D}$, where $\mathcal{D}$ is defined similarly to $\mathcal{R}$.

The following lemma shows that the dynamics of the tracking error $e_t \doteq r_t - x_t$ can be described by a quasi-LPV (Linear Parameter Varying) equation.

Lemma 1. *Consider the closed-loop system (5). The dynamics of the tracking error $e_t \doteq r_t - x_t$ is described by*

$$e_{t+1} = A_t e_t + \xi_t - d_t \tag{6}$$

where $\xi_t \doteq r_{t+1} - f_c(r_t, \mathbf{r}_t)$ and

$$A_t \doteq \begin{cases} \frac{(f_c(r_t,\mathbf{r}_t) - f_c(x_t,\mathbf{r}_t))(r_t - x_t)^\top}{\|r_t - x_t\|^2}, \ x_t \neq r_t \\ \mathbf{0}, \ x_t = r_t. \end{cases} \tag{7}$$

Proof. The proof is provided in [15]. □

Besides providing a quasi-LPV representation of the tracking error dynamics, the lemma shows that this dynamics is driven by two external variables: the disturbance d_t and the quantity ξ_t. The latter can be interpreted as a "reachability" variable, measuring the "smoothness" of the reference signal. Indeed, suppose for simplicity that the matrix R in (2) is null. Consider the ideal situation where, at a time t, $x_t = r_t$ and a c exists such that $r_{t+1} = f(r_t, S_1^u \Gamma c^*(x_t, \mathbf{r}_t))$. Then, minimizing the objective function yields $\xi_t = 0$. In such an ideal situation, ξ_t is null and the tracking error dynamics is only driven by the disturbance. In generic situations, "smooth" reference signals are expected to yield small values of ξ_t, while references with abrupt variations give large values of ξ_t.

Based on the above lemma, a finite-gain stability result is now derived, using the following assumption.

Assumption 1. *Closed-loop transition matrix bound. Define the transition matrix*

$$F_t^\tau \doteq \begin{cases} A_{t+\tau-1} A_{t+\tau-2} \dots A_t, \ \tau > 0 \\ I, \ \tau = 0 \end{cases} \tag{8}$$

where A_i is given in (7). We assume that an integer $\tau > 0$ exists such that

$$\gamma \equiv \gamma(\tau) \doteq \max_{l \geq 0, x_0 \in X_0, \boldsymbol{r} \in \mathcal{R}, \mathbf{d} \in \mathcal{D}} \|F_l^\tau\| < 1. \quad \square \tag{9}$$

This assumption is based on the quasi-LPV tracking error dynamics (6), whose properties can be studied using the theory of Linear Time Varying (LTV) systems. Most stability assumptions that can be found in the LTV literature regard the matrix A_l. For example, a standard assumption is that A_l has eigenvalues inside the unit circle and is "slowly varying" [16]. Another typical assumption is $\|A_l\| < 1$, $l = 0, 1, \dots$, [16]. In our approach, we do not make assumptions on A_l but on the matrix F_l^τ. This latter is the product of several subsequent matrices A_l, making our assumption weaker than the standard ones regarding directly A_l. As shown by the following theorem, (9) is the key condition for guaranteeing closed-loop finite-gain stability.

Theorem 1. *Let Assumption 1 hold. Then:*

(i) The closed-loop system (5) is finite-gain stable on $(X_0, \mathcal{R})$.

(ii) The tracking error $e_t \doteq r_t - x_t$ is bounded as

$$\begin{aligned} \|e_t\| \leq & \Lambda_0 \rho^t \|e_0\| + \Lambda_\xi \sum_{i=0}^{t-1} \rho^{t-i} \|\xi_i\| \\ & + \Lambda_d \sum_{i=0}^{t-1} \rho^{t-i} \|d_i\| \end{aligned} \tag{10}$$

where $\rho \doteq \gamma^{1/\tau} < 1$, and Λ_0, Λ_ξ, Λ_d are finite non-negative constants.

Proof. The proof is provided in [15]. □

Remark 1. *According to Theorem 1, the key condition for closed-loop finite-gain stability is $\gamma < 1$, see (9). As discussed above, this condition relies on a quasi-LPV representation of the tracking error dynamics and is less conservative than the standard assumptions that can be used for quasi-LPV/LTV systems. An important observation is that, while checking FGS using Definition 1 is practically impossible, the verification of (9) is possible. Indeed, in Section IV we propose a viable method for the numerical estimation of γ, allowing us to enforce stability in practice. The estimate of γ is also the core of a systematic procedure for efficient NMPC design, see Sections IV and V.* □

IV. NMPC DESIGN

This section presents a systematic NMPC design methodology, allowing an effective selection of a suitable NMPC configuration leading to closed-loop finite-gain stability, as defined by Theorem 1, while providing satisfactory tracking performance. The design procedure consists of three key steps: (i) determination of the integer horizon τ used in the finite-gain analysis; (ii) a closed-loop function sampling campaign that generates candidate NMPC configurations and closed-loop data over a finite interval of length τ; (iii) a selection rule that chooses the best configuration according to both tracking error and stability criteria.

Let Φ denote a finite set of NMPC configurations. Each configuration $\phi \in \Phi$ is described by a prediction horizon T and weighting matrices $Q \succ 0$, $R \succ 0$, $P \succ 0$. Let $s = (x_0, \boldsymbol{r}, \boldsymbol{d})$ be a generic scenario containing an initial state $x_0 \in X_0$, a reference sequence $\boldsymbol{r} = (r_0, r_1, \dots) \in \mathcal{R}$ and an exogenous disturbance sequence $\boldsymbol{d} = (d_0, d_1, \dots) \in \mathcal{D}$.

At a given time t and for a fixed configuration ϕ and scenario s, denote by: i) $x_t(\phi, s)$ the closed-loop state trajectory produced by the NMPC law; ii) $e_t(\phi, s)$ the corresponding tracking error; and iii) $F_t^\tau(\phi, s)$ the transition matrix.

It is important to highlight that the proposed procedure does not provide a formal guarantee of finite-gain stability

for all initial conditions $x_0 \in X_0$. Instead, it offers a practical approach to identify NMPC configurations that are likely to satisfy the finite-gain condition over a predefined operating domain, described through a finite set of representative scenarios. The reliability of the method therefore depends on the coverage of the scenario sets and on the robustness margins introduced in the design.

A. Selection of the simulation horizon τ

The integer τ, introduced in Assumption 1, must be chosen sufficiently large to ensure that the norm of the product $F_t^{\tau}(\phi, s)$, denoted by $\gamma(\tau)$, remains strictly less than one for all the considered scenarios. Since the exact verification of this condition over the full domain is generally intractable, we estimate a conservative value of τ by means of the following procedure:

1) Generate an initial scenario set $\mathcal{S}_{\text{in}} = \{s^1, \dots, s^M\}$, where each element is defined as $s^i = (x_0^i, \boldsymbol{r}^i, \boldsymbol{d}^i)$.
2) Choose a nominal NMPC configuration ϕ_{nom}, denoted by the tuple $\phi_{\text{nom}} = (T_{\text{nom}}, Q_{\text{nom}}, R_{\text{nom}}, P_{\text{nom}})$, obtained, for instance, from standard design heuristics such as Bryson's rule [17].
3) For each scenario $s^i \in \mathcal{S}_{\text{in}}$, perform a closed-loop simulation, using ϕ_{nom}, until a final time T_{fin} (e.g., the time at which the system reaches the reference or a predefined maximum simulation time). At each time $t = 1, \dots, T_{\text{fin}}$ compute the value of γ as defined in Eq. (9) and determine the smallest integer τ_i such that
$$\gamma_i(t) < 1, \quad \forall t = \tau_i, ...T_{\text{fin}}. \tag{11}$$
4) Compute τ^* as
$$\tau^\star = \max_{i \in \{1,\dots,M\}} \tau_i + \Delta_\tau, \tag{12}$$
where $\Delta_\tau \geq 0$ is a robustness margin.

The role of Δ_τ is to improve the reliability of the finite-gain condition with respect to untested scenarios. Since $\tau^\star$ is computed from a finite scenario set, it may underestimate the true contraction horizon required over the entire set X_0. The addition of Δ_τ compensates for this effect and increases the likelihood that the finite-gain condition $\gamma(\tau^\star) < 1$ also holds for scenarios not explicitly included in $\mathcal{S}_{\text{in}}$. However, this mechanism does not provide a formal guarantee over all $x_0 \in X_0$, and its effectiveness depends on the representativeness of the scenario set.

If (12) cannot be satisfied for any τ, it may be necessary to modify ϕ_{nom} (e.g. increase T or retune Q, R, P) or relax performance requirements.

B. Closed-loop Function Sampling

A Closed-loop Function Sampling campaign is conducted to evaluate the finite-time closed-loop response for each NMPC configuration. The steps are:

1) Generate an evaluation scenario set $\mathcal{S} = \{s^1, \dots, s^N\}$, with $N \geq M$. This set is typically constructed in order to ensure a proper coverage of the relevant initial states, reference trajectories, and disturbance profiles, and may be an extension of the initial scenario set, i.e., $\mathcal{S} \supseteq \mathcal{S}_{\text{in}}$.
2) Build the NMPC configuration set $\Phi = \{\phi_j\}_{j=1}^K$, $\phi_j = (T^{(j)}, Q^{(j)}, R^{(j)}, P^{(j)})$, by choosing different values of T, Q, R and P. The matrices Q, R and P can be scaled starting from the nominal NMPC configuration $Q = \alpha_Q Q_{\text{nom}}$, $R = \alpha_R R_{\text{nom}}$, $P = \alpha_P P_{\text{nom}}$ where $\alpha_Q, \alpha_R, \alpha_P > 0$ are scalar tuning parameters.
3) For each configuration ϕ_j, simulate the closed-loop state trajectory $x_t(\phi_j, s^i)$ over the time interval $t = 1, \dots, \tau^*$ for all scenarios s^i, and store the resulting tracking error $e_{\tau^*}(\phi_j, s^i)$ and transition matrix $F_t^{\tau^*}(\phi_j, s^i)$.

C. Selection of the optimal NMPC configuration

For each NMPC configuration $\phi \in \Phi$, we define two metrics:

a) Tracking error metric: The worst-case tracking error over all scenarios is defined as
$$E(\phi) = \max_{i=1,\dots,N} E(\phi, s^i) = \max_{i=1,\dots,N} \|e_{\tau^*}(\phi, s^i)\|. \tag{13}$$

b) Finite-gain stability metric: The worst-case finite-gain stability index is given by
$$L(\phi) = \max_{i=1,\dots,N} \gamma_i(\tau^*) = \max_{i=1,\dots,N} \|\big(F_t^{\tau^*}(\phi, s^i)\big)\|. \tag{14}$$

For the selection rule, we first normalize the two metrics relative to their worst-case values across the entire set Φ
$$E_n(\phi) = \frac{E(\phi)}{\max_\Phi E}, \qquad L_n(\phi) = \frac{L(\phi)}{\max_\Phi L},$$
and then define the following cost function
$$J_{sr}(\phi) = \alpha_J E_n(\phi) + (1-\alpha_J) L_n(\phi), \qquad \alpha_J \in [0,1]. \tag{15}$$
Finally, the optimal NMPC configuration is selected as
$$\phi^\star = \arg\min_{\phi \in \Phi} J_{sr}(\phi). \tag{16}$$

D. Visualization and sensitivity analysis

To improve interpretability of the results, the sensitivity of the finite-gain stability index (and tracking performance) to key design parameters can be analyzed through parametric analysis. Specifically, surface plots or level curves of $L(\phi)$ and $E(\phi)$ are generated over a two-dimensional grid defined by the prediction horizon T and the ratio $\kappa = \alpha_Q/\alpha_R$. These plots show how variations in T and in the relative weighting between states and inputs influence the closed-loop stability margin (and tracking accuracy).

E. Summary of the design algorithm

Algorithm 1 summarizes the proposed procedure.

By exploiting the finite-gain stability condition of Theorem 1, the proposed procedure provides a systematic and computationally efficient method for obtaining an optimal NMPC configuration through a function sampling campaign in which each closed-loop simulation is performed only

Algorithm 1 Selection of the optimal NMPC configuration via finite-gain analysis

Input: System (1), X_0, $\mathcal{R}$, $\mathcal{D}$
Output: $\phi^\star$

```
Choose nominal φ_nom and compute τ* according to (12).
Construct the NMPC configuration set Φ and the scenario set S.
for all φ ∈ Φ do:
    for all s^i ∈ S do:
        Simulate closed-loop for t = 1, ..., τ*.
        Compute E(φ, s^i) and F_t^{τ*}(φ, s^i).
    end for
    Compute E(φ) and L(φ).
end for
Select φ* through Eq. (16).
Produce surface plots and, if necessary, refine the candidate grid.
```

over the horizon τ^*. This significantly reduces the computational effort compared to traditional long-horizon simulations, where different types of scenarios are considered and, for each of them, several maneuvers are simulated for their complete duration.

V. Example: Lateral/longitudinal control for an automated vehicle

We present the application of the proposed NMPC design methodology to the problem of controlling the lateral and longitudinal dynamics of an automated vehicle. The resulting controller is then validated on two standard maneuvers: lane keeping and obstacle avoidance.

A standard Dynamic Single-Track (DST) model is considered, providing a representation of both lateral and longitudinal vehicle dynamics:

$$\begin{aligned}
\dot{p}_X &= v_x \cos\psi - v_y \sin\psi \\
\dot{p}_Y &= v_x \sin\psi + v_y \cos\psi \\
\dot{\psi} &= \omega \\
\dot{v}_x &= v_y \dot{\psi} + a_x \\
\dot{v}_y &= -v_x \dot{\psi} + \frac{2}{m}\left(F_{yf} + F_{yr}\right) \\
\dot{\omega} &= \frac{2}{I_z}\left(l_f F_{yf} - l_r F_{yr}\right)
\end{aligned} \tag{17}$$

where p_X and p_Y are the coordinates of the vehicle in an inertial frame, ψ is the yaw angle, ω denotes the yaw rate, and v_x, v_y are the longitudinal and lateral speeds, respectively. The main physical parameters of the model and their values are as follows: mass $m = 1575\,\text{kg}$; moment of inertia $I_z = 4000\,\text{kgm}^2$; distances of Center of Gravity (CoG) to front axle $l_f = 1.2\,\text{m}$ and rear axle $l_r = 1.6\,\text{m}$. Moreover, F_{yf} and F_{yr} are the lateral forces, given by $F_{yf} = -c_f\beta_f$, $F_{yr} = -c_r\beta_r$, where $c_f = 2.7 \cdot 10^4\,\text{N/rad}$ and $c_r = 2 \cdot 10^4\,\text{N/rad}$ are the front/rear cornering stiffnesses. The front/rear tire slip angles β_f and β_r are defined as

$$\beta_f = \operatorname{atan}\left(\frac{v_y + l_f\dot{\psi}}{v_x}\right) - \delta_f,\ \beta_r = \operatorname{atan}\left(\frac{v_y - l_r\dot{\psi}}{v_x}\right). \tag{18}$$

The longitudinal acceleration a_x and the steering angle δ_f are the control variables with the following saturation constraints: $a_x \in [-5, 3]\ \text{m/s}^2, \delta_f \in [-0.78, 0.78]\ \text{rad}$. The DST equations have been discretized by means of the Forward Euler method, using the sampling time $Ts = 0.1$ s. The controlled outputs considered in the NMPC design are (p_X, p_Y, ψ, v_x). To account for model–plant mismatch and measurement errors, a Gaussian noise term was added to the plant dynamics during simulation. Specifically, the noise had a mean vector $\mu = [0.5, 0.5, 0.001, 0.05, 0.05, 0.001]$ and covariance matrix $\Sigma = \text{diag}(10^{-4}, 10^{-4}, 10^{-5}, 10^{-4}, 10^{-5}, 10^{-5})$.

The NMPC configuration is obtained by following the procedure detailed in Section IV. The scenario sets $\mathcal{S}_{\text{in}}$ and $\mathcal{S}$ are generated using the Latin Hypercube Sampling (LHS) [18] which produces quasi-random samples of parameter values from a multidimensional distribution, providing better coverage of the domain compared to random sampling. The initial state set X_0 is defined as:

$$\begin{aligned}
X_0 = \{x_0 = (p_{X,0}, p_{Y,0}, \psi_0, v_{x,0}, v_{y,0}, \omega_0) \mid \\
p_{X,0} \in [0, 100]\ \text{m},\ p_{Y,0} \in [0, 100]\ \text{m},\ \psi_0 \in [-0.5, 0.5]\ \text{rad}, \\
v_{x,0} \in [5, 15]\ \text{m/s},\ v_{y,0} \in [0.1, 1]\ \text{m/s}, \\
\omega_0 \in [-0.1, 0.1]\ \text{rad/s}\}.
\end{aligned}$$

The reference set $\mathcal{R}$ was generated using the same intervals for the corresponding reference variables (position, yaw angle and longitudinal velocity), ensuring that initial conditions and targets are sampled from a consistent domain.

1) Selection of the simulation horizon τ^*: For each scenario $s^i \in \mathcal{S}_{\text{in}}$, closed-loop simulations were performed until the vehicle reached the reference trajectory. The simulation horizon was set to $\tau^* = 15$.

2) Closed-loop Function Sampling campaign: The sampling campaign was conducted by evaluating a set of $K = 160$ candidate NMPC configurations. For each configuration ϕ_j, closed-loop simulations were performed over the finite simulation horizon τ^* and for all scenarios in S, resulting in more than 3000 simulations. The complete design procedure took about 40 minutes, using a laptop with processor Intel(R) Core(TM) i7-10750H CPU @ 2.60GHz.

3) Selection of the optimal NMPC configuration: After computing the normalized indices $E_n(\phi)$ and $L_n(\phi)$, the scalar cost function $J_{sr}(\phi) = \alpha_J E_n(\phi) + (1 - \alpha_J) L_n(\phi)$ was minimized with $\alpha_J = 0.5$, providing a good balance between tracking accuracy and stability.

4) Visualization and sensitivity analysis: To illustrate the influence of design parameters, level curve plots of the metrics $E(\phi)$ and $L(\phi)$ were generated over the grid (T, κ) (see Figure 1).

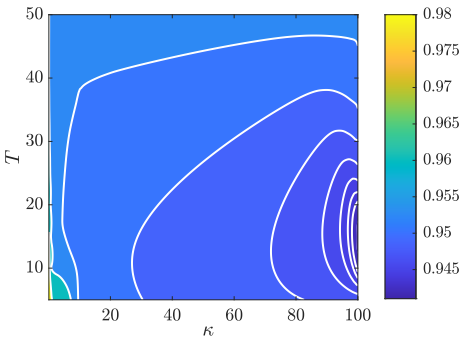


Fig. 1. Level curves of $L(\phi)$.

The selected NMPC configuration $\phi^\star$ was first validated on a lane-keeping scenario and compared with an alternative configuration ϕ_a. Both controllers were tested on a curvilinear urban road, modeled by a sinusoidal path of the form $r = 8\sin(0.02\,t)$. The longitudinal velocity reference was set to a constant value of $v_{x,\mathrm{ref}} = 40$ km/h. Table I summarizes the key tracking error metrics. The results indicate that $\phi^\star$ achieves marginally better lateral tracking and smaller orientation deviations than ϕ_a.

TABLE I

COMPARISON BETWEEN $\phi^\star$ AND ϕ_a.

	$\phi^\star$	ϕ_a
RMS lateral error [m]	0.025	0.029
Maximum lateral error [m]	0.07	0.1
RMS orientation error [rad]	0.0005	0.00055
Maximum orientation error [rad]	0.0007	0.0012

The second validation scenario involves an obstacle-avoidance maneuver. The obstacle is modeled as an elliptical safety region with center coordinates $(c_X, c_Y) = (5, 5)$ and sizes $(l_X, l_Y) = (1, 1)$. As shown in Figure 2, the optimal configuration $\phi^\star$ maintained the vehicle trajectory closer to the nominal path connecting the initial and final points, thus minimizing lateral deviation during the maneuver.

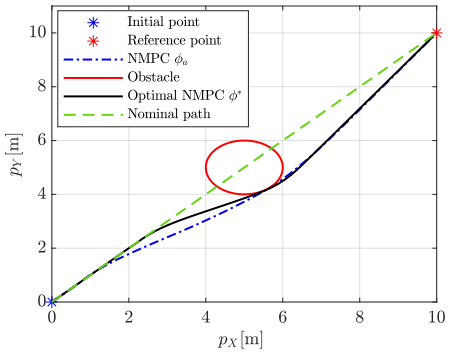


Fig. 2. Obstacle avoidance scenario.

VI. CONCLUSIONS

This paper has presented a finite-gain-based approach to the stabilizing design of NMPC control systems. The main result is a bound on the closed-loop tracking error under a finite-step contraction condition on the trajectory-dependent error dynamics, which makes explicit the role of the initial error, the disturbance, and the reference variation. An offline scenario-based design procedure has been proposed to select the prediction horizon and weighting matrices by combining a tracking-performance index with a finite-gain stability index computed over a finite simulation horizon. The method has been validated on an automated-vehicle example, where it identified a controller configuration with satisfactory tracking and stability properties over the vehicle maneuvers.

REFERENCES

[1] J. Rawlings, D. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, 2017.
[2] D. Mayne, "An apologia for stabilising terminal conditions in model predictive control," *Int. J. Control*, vol. 86, no. 11, pp. 2090 – 2095, 2013.
[3] A. R. Teel, *Optimal control, stabilization and nonsmooth analysis*. Springer, 2004, vol. 301, ch. Discrete time receding horizon control: is the stability robust, pp. 3 – 28.
[4] C. M. Kellet and A. R. Teel, "Discrete-time asymptotic controllability implies smooth control-lyapunov function," *Sys. Cont. Let.*, vol. 52, pp. 349 – 359, 2004.
[5] D. Limon, T. Alamo, D. M. Raimondo, D. M. de la Peña, J. M. Bravo, A. Ferramosca, and E. F. Camacho, *Input-to-State Stability: A Unifying Framework for Robust Model Predictive Control*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 1–26.
[6] L. Zaccarian, A. R. Teel, and D. Nešić, "On finite gain $\mathcal{L}_p$ stability of nonlinear sampled-data systems," *Sys. Cont. Let.*, vol. 49, no. 3, pp. 201–212, 2003.
[7] H. K. Khalil, *Nonlinear Systems*. Upper Saddle River, New Jersey, USA: Prentice Hall, 2002.
[8] D. Limon, T. Alamo, F. Salas, and E. F. Camacho, "Input to state stability of min–max MPC controllers for nonlinear systems with bounded uncertainties," *Automatica*, vol. 42, no. 5, pp. 797–803, 2006.
[9] D. Limon, T. Alamo, and E. Camacho, "Input-to-state stable MPC for constrained discrete-time nonlinear systems with bounded additive uncertainties," in *41st IEEE CDC, 2002*, vol. 4, 2002, pp. 4619–4624.
[10] J. Köhler, M. A. Müller, and F. Allgöwer, "Analysis and design of model predictive control frameworks for dynamic operation—an overview," *Annu. Rev. Control*, vol. 57, p. 100929, 2024.
[11] P. Krupa, J. Köhler, A. Ferramosca, I. Alvarado, M. Zeilinger, T. Alamo, and D. Limon, "Model predictive control for tracking using artificial references: Fundamentals, recent results and practical implementation," in *IEEE 63rd CDC*, 2024, pp. 2977–2991.
[12] M. Pagone, M. Boggio, C. Novara, A. Proskurnikov, and G. C. Calafiore, "Continuous-time nonlinear model predictive control based on Pontryagin minimum principle and penalty functions," *Int. J. Control*, vol. 98, no. 4, p. 765–781, 2025.
[13] D. Limon, I. Alvarado, T. Alamo, and E. F. Camacho, "MPC for tracking piecewise constant references for constrained linear systems," *Automatica*, vol. 44, no. 9, pp. 2382–2387, 2008.
[14] R. C. Shekhar and J. M. Maciejowski, "Robust variable horizon MPC with move blocking," *Sys. Cont. Let.*, vol. 61, no. 4, pp. 587–594, 2012.
[15] C. Novara, M. Boggio, L. Calogero, and M. Pagone, "A finite-gain stability approach to NMPC design: the extended version," *ArXiv, https://arxiv.org/abs/2606.18863*, 2026.
[16] W. J. Rugh, *Linear System Theory*. Prentice Hall, 1996.
[17] A. E. Bryson and Y. C. Ho, *Applied Optimal Control*. Blaisdell Publishing Company, 1969.
[18] M. D. McKay, R. J. Beckman, and W. J. Conover, "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," *Technometrics*, vol. 21, no. 2, pp. 239–245, 1979.