On the Regret of Model Predictive Control With Imperfect Inputs

Liu, Changrui; Shi, Shengling; Schutter, Bart De

DOI
10.1109/LCSYS.2025.3577083

Publication date
2025

Document Version
Final published version

Published in
IEEE Control Systems Letters

On the Regret of Model Predictive Control With Imperfect Inputs

Changrui Liu, Shengling Shi, and Bart De Schutter, *Fellow, IEEE*

***Abstract*—Implementing model predictive control (MPC) in practice faces many subtle but prevalent problems, including modeling errors, solver errors, and actuator faults. In essence, the real control input applied to the system always deviates from the ideal one based on a perfect controller, resulting in an imperfect controller. In this letter, we provide a general analysis to quantify the suboptimality of MPC for Lipschitz-continuous nonlinear systems due to imperfect control inputs in terms of *dynamic regret*. Based on a general assumption about how the imperfect controller may improve over time, sublinear regret upper bounds are established for cases where the closed-loop system under the ideal controller is Lipschitz-contractive (i.e., its Lipschitz constant is smaller than one). In addition, we also discuss how the regret scales when the closed-loop system under the oracle controller is not Lipschitz-contractive. The results provide insights into designing suitable MPC strategies, especially for learning-based MPC.**

***Index Terms*—Optimal control, predictive control, regret analysis, input errors.**

I. Introduction

MODEL Predictive Control (MPC) is a powerful control method due to its ability to account for future behavior, handle constraints, and provide closed-loop guarantees (e.g., stability and robustness) [1]. The advantages of MPC have facilitated numerous applications in fields like robotics [2] and aerospace [3]. However, despite its theoretical soundness, implementing MPC in real-world scenarios often encounters limitations that compromise its performance. Common problems include model mismatch [4], [5], solver errors [6], and actuator faults [3]. Although the source of these problems differs, the common result is having discrepancies between the ideal input derived with all ideal assumptions (e.g., the model is fully known, and the actuator is flawless) and the actual input applied in practice. Notably, the imperfect input causes accumulated deviations from the ideal trajectory, leading to a gap between the expected and the actual control performance.

To improve the performance of MPC due to previously mentioned factors, several strategies have been developed, e.g., fault-tolerant MPC [3], adaptive MPC [7], and learning-based MPC [8]. Nevertheless, obtaining the true model, an exact optimizer, and accurate fault estimation is demanding. Therefore, imperfect inputs are unavoidable, and analyzing the performance gap due to the associated input deviation is thus crucial for developing improved MPC strategies when perfect inputs cannot be realized. This letter focuses on theoretical performance analysis of MPC in the presence of imperfect inputs instead of providing specific MPC designs.

In terms of the performance analysis of MPC, suboptimality due to model mismatch has been studied, focusing on both the transient performance using *dynamic regret* [2], [8], [9] as well as infinite-horizon performance [4], [5]. Besides, several methods have been developed to reduce suboptimality caused by modeling errors, e.g., Bayesian learning with active exploration [8], recursive-least-square (RLS) identification with deterministic [7] or probabilistic guarantees [2], and offset-free MPC design for improved tracking performance [10], [11]. The suboptimality due to early termination of the solver has also been studied [6]. For actuator faults, the performance analysis of MPC is an open problem. Moreover, the existing results are mostly for linear systems [4], [5], [6] or assume a dominant linear model [8]. For nonlinear systems, the effect of parametric modeling errors has been investigated [9], and the suboptimality due to a general imperfect controller has been studied in [12] under the assumption that the closed-loop system under the ideal controller satisfies a type of incremental stability condition. Considering that all of the contributions towards the performance analysis of MPC in the literature studied a tailored MPC controller (e.g., MPC with Bayesian learning [8] and certainty-equivalence MPC [4], [9]) and linear models are still dominantly considered, it is imperative to focus on a more general class of systems and to develop a more fundamental analysis that can benefit the performance analysis due to general imperfect inputs.

Contributions: In this letter, a general analysis is provided to quantify the suboptimality of MPC in terms of *dynamic regret* when imperfect control inputs are used. When the ideal MPC is employed, the resulting closed-loop system is called the *oracle* closed-loop system. Compared to most of the existing results on linear models [4], [6], [8], we focus on Lipschitz-continuous nonlinear systems with also Lipschitz-continuous cost functions. More importantly, our analysis framework does not specify the error source of the control input or the specific MPC controller, encompassing

Received 11 March 2025; revised 10 May 2025; accepted 25 May 2025. Date of publication 5 June 2025; date of current version 19 June 2025. This work was supported by the European Research Council (ERC) under the European Union's Horizon 2020 Research and Innovation Programme under Grant 101018826 - CLariNet. Recommended by Senior Editor S. Olaru. *(Corresponding author: Changrui Liu.)*

Changrui Liu and Bart De Schutter are with the Delft Center for Systems and Control, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: C.Liu-14@tudelft.nl; B.Deschutter@tudelft.nl).

Shengling Shi is with the Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: slshi@mit.edu).

Digital Object Identifier 10.1109/LCSYS.2025.3577083

a broad range of MPC control problems. The most relevant work to our results is [12], compared to which the major differences of our results are that i) only the simpler and less restrictive Lipschitzness is assumed to quantify the trajectory perturbation instead of using the E-δ-ISS property or other related stability notions [13], ii) sufficient conditions on the online learning (adaptation) rate of the suboptimal MPC controller is specified to guarantee sublinear regret when the oracle closed-loop system is Lipschitz-*contractive* (i.e., its Lipschitz constant is smaller than 1), and iii) the cases where the oracle closed-loop system is Lipschitz-*noncontractive* (i.e., its Lipschitz constant is greater than or equal to 1) are also discussed. The regret analysis in this letter is the first in the literature to preserve this high-level generality, making the results applicable in different scenarios. Specifically, given some additional assumptions on the control laws, we derive regret upper bounds and provide sufficient conditions ensuring *sublinear* regret when the closed-loop system under the ideal control inputs is Lipschitz-contractive.

The remainder of this letter is organized as follows. In Section II, the MPC control problem is introduced. Section III formulates the performance analysis problem, where the regret metric is defined. The analysis pipeline is provided in Section IV, and Section V concludes this letter.

II. Preliminaries

A. Notation

The sets of real and non-negative real numbers are denoted by $\mathbb{R}$ and $\mathbb{R}_+$, respectively. The set of natural numbers is $\mathbb{N}$, and $\mathbb{I}_{[a,b]} := \mathbb{N} \cap [a, b]$. The two-norm is denoted by $\|\cdot\|$. The sum $\sum_{i=j_1}^{j_2}$ and product $\prod_{i=j_1}^{j_2}$ with $j_2 < j_1$ are defined, respectively, as 0 and 1. A function $\alpha : [0, a) \to \mathbb{R}_+$ for some $a \in \mathbb{R}_+$ is said to belong to class $\mathcal{K}$ (i.e., $\alpha \in \mathcal{K}$) if $\alpha(0) = 0$ and α is strictly increasing. Asymptotic bounds using Big-O notation are obtained under $T \to \infty$ with T being the horizon.

B. Optimal Control Problem

Consider a discrete-time dynamical system given by

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

where $x_t \in \mathcal{X} = \mathbb{R}^n$ and $u_t \in \mathcal{U} \subseteq \mathbb{R}^m$ are, respectively, the state and input at time step t, and the function $f : \mathcal{X} \times \mathcal{U} \to \mathcal{X}$ is Lipschitz continuous in the state x and input u, i.e., $\forall x', x'' \in \mathcal{X}$ and $\forall u', u'' \in \mathcal{U}$, it holds that

$$\|f(x', u') - f(x'', u'')\| \leq L_{f,x}\|x' - x''\| + L_{f,u}\|u' - u''\|, \tag{2}$$

where $L_{f,x}$ and $L_{f,u}$ are the Lipschitz constants. Input affine nonlinear systems satisfy (2) if $\mathcal{X}$ is bounded [12], and similar assumptions are also frequently used in MPC [8], [14]. The set $\mathcal{U}$ is assumed to be *compact*, and $\mathcal{U} := \{u \in \mathbb{R}^m \mid g_u(u) \leq \mathbf{0}\}$, where $g_u : \mathbb{R}^m \to \mathbb{R}^{c_m}$. We further define $\mathrm{diam}(\mathcal{U}) := \sup_{u_1,u_2 \in \mathcal{U}} \|u_1 - u_2\|$, where $\mathrm{diam}(\mathcal{U}) < +\infty$ as $\mathcal{U}$ is bounded. At time step t, the system incurs a stage cost $\ell(x_t, u_t)$, where the function $\ell : \mathcal{X} \times \mathcal{U} \to \mathbb{R}$ is also Lipschitz continuous, i.e., $\forall x', x'' \in \mathcal{X}$ and $\forall u', u'' \in \mathcal{U}$, it holds that

$$|\ell(x', u') - \ell(x'', u'')| \leq L_{\ell,x}\|x' - x''\| + L_{\ell,u}\|u' - u''\|, \tag{3}$$

where $L_{\ell,x}$ and $L_{\ell,u}$ are the Lipschitz constants. Lipschitz-continous costs are commonly used in MPC [8], [9] and optimal control [12], and examples are the linear cost in economic MPC [15] and 1-norm or ∞-norm costs [16]. Given a prediction horizon $N \in \mathbb{I}_{[1:\infty]}$, at each time step t, the MPC controller solves the following optimization problem:

$$\begin{aligned}
\mathrm{P}_{\mathrm{MPC}}(x_t): \quad & \min_{\{u_{k|t}\}_{k=0}^{N-1}, \{x_{k|t}\}_{k=0}^{N}} F(x_{N|t}) + \sum_{k=0}^{N-1} \ell(x_{k|t}, u_{k|t}) \\
\text{s.t.} \quad & x_{k+1|t} = f(x_{k|t}, u_{k|t}), \forall k \in \mathbb{I}_{[0:N-1]}, \\
& g_u(u_{k|t}) \leq \mathbf{0}, \forall k \in \mathbb{I}_{[0:N-1]}, \\
& x_{0|t} = x_t,
\end{aligned}$$

where $x_{k|t}$ and $u_{k|t}$ are the k-step-ahead predicted state and input at time step t, $F : \mathcal{X} \to \mathbb{R}$ is the *terminal cost* that is used to approximate the tail cost for the associated infinite-horizon optimal control problem [1], [16], and x_t is the measured state.[1] Solving $\mathrm{P}_{\mathrm{MPC}}$ returns the inputs $\{u^\star_{k|t}(x_t)\}_{k=0}^{N-1}$, and only $u^\star_{0|t}(x_t)$ will be applied. This optimization problem defines an implicit control law as $\mu_{\mathtt{OPT}}(x_t) := u^\star_{0|t}(x_t)$ [1]. The subscript $\mathtt{OPT}$ indicates that $\mu_{\mathtt{OPT}} : \mathcal{X} \to \mathcal{U}$ is the ideal control law, under which the closed-loop system is

$$x^\star_{t+1} = f_{\mathtt{OPT}}(x^\star_t) := f(x^\star_t, \mu_{\mathtt{OPT}}(x^\star_t)), \tag{4}$$

where $x^\star_t$ is the closed-loop state generated by (4) under the ideal optimal control law $\mu_{\mathtt{OPT}}$ at time step t. In the sequel, $\mathtt{OPT}$ will be referred to as the *oracle* controller that has perfect knowledge of the true model and can calculate the optimal input without numerical errors.

III. Problem Formulation

The ideal control law is based on perfect control, neglecting practical considerations. The actual control applied to the system is always a *perturbed* one, and a practical controller could be designed that may have more functionalities (e.g., model learning [8], [17] or fault detection and isolation [3]), or that may follow the baseline design without any additional modules to tackle those practical problems (e.g., certainty-equivalence MPC [4], [5], [9]).

We preserve the generality in this letter and uniformly indicate the practical MPC controller by $\mathtt{ALG}$ and its induced (time-varying) control law as $\mu_{\mathtt{ALG},t}$. Then, the (time-varying) closed-loop system under $\mu_{\mathtt{ALG},t}$ is given by

$$\tilde{x}_{t+1} = f_{\mathtt{ALG},t}(\tilde{x}_t) := f(\tilde{x}_t, \mu_{\mathtt{ALG},t}(\tilde{x}_t)), \tag{5}$$

where $\tilde{x}_t$ denotes the state evolving under the controller $\mathtt{ALG}$.

Assumption 1 (Existence of Invariant Set): There exists a *compact* set $\mathcal{X}_{\mathrm{inv}} \subset \mathcal{X}$ such that it is *positive invariant* under $f_{\mathtt{OPT}}$ and $f_{\mathtt{ALG},t}$, $(t = 0, 1, 2, \dots)$, i.e., if $x \in \mathcal{X}_{\mathrm{inv}}$, it holds that $f_{\mathtt{OPT}}(x) \in \mathcal{X}_{\mathrm{inv}}$ and $\forall t \in \mathbb{N}$, $f_{\mathtt{ALG},t}(x) \in \mathcal{X}_{\mathrm{inv}}$.

Assumption 1 states that there exists an implicit state constraint set that will not be violated for the considered input-constrained control problem, which eases the analysis of state perturbation. Similar assumptions have been made in [12]. Moreover, since the set $\mathcal{X}_{\mathrm{inv}}$ is compact and thus bounded, the Lipschitz-continuous stage cost requirement (cf. (3)) becomes less restrictive since the popular quadratic function $\ell(x, u) = x^\top Qx + u^\top Ru$ also fits the framework.

[1] In this letter, state-feedback MPC is considered and it is assumed that the state can be accurately measured.

To evaluate the performance of ALG compared to OPT, the *dynamic regret* is used as the metric. Starting from $x = x_0^\star = \tilde{x}_0 \in \mathcal{X}_{\text{inv}}$, the dynamic regret Reg_T is defined as

$$\text{Reg}_T := \sum_{t=0}^{T-1} \left[\ell\big(\tilde{x}_t, \mu_{\texttt{ALG},t}(\tilde{x}_t)\big) - \ell\big(x_t^\star, \mu_{\texttt{OPT}}(x_t^\star)\big)\right]. \quad (6)$$

The problem is to find a theoretical framework to derive an upper bound of Reg_T and to specify conditions under which the regret is *sublinear* in T, i.e., $\lim_{T\to\infty} \text{Reg}_T/T = 0$. A sublinear regret upper bound implies that ALG performs at least as well as OPT in the long run. It is also noted that Reg_T is a function of the prediction horizon N. However, this explicit dependency on N is out of the scope of the analysis pipeline in this letter and is left for future work.

Remark 1: One important aspect of MPC controllers is stability, which is often guaranteed by design or imposing additional assumptions. Stability is not needed in the analysis of this letter, and the primary concern of the MPC controller formulated as in $\text{P}_{\text{MPC}}(x_t)$ is to minimize the cost in the long run instead of regulating the state to a certain equilibrium. However, as will be discussed in Remark 4, stability can help achieve sublinear regret under weaker conditions.

IV. Regret Analysis

In this section, the regret analysis is provided. Section IV-A presents assumptions and preparatory analysis. In Section IV-B, a refined analysis under the contractive perturbation condition is given, extending the results of [12]. The final Section IV-C discusses the case without contractivity.

A. Preparatory Analysis

Based on the Lipschitz continuity of ℓ, the regret can be bounded as

$$\text{Reg}_T \le \sum_{t=0}^{T-1} \left[L_{\ell,x}\|\tilde{x}_t - x_t^\star\| + L_{\ell,u}\|\mu_{\texttt{ALG},t}(\tilde{x}_t) - \mu_{\texttt{OPT}}\big(x_t^\star\big)\|\right], \quad (7)$$

where $\mu_{\texttt{ALG},t}(\tilde{x}_t) - \mu_{\texttt{OPT}}(x_t^\star)$ needs to be evaluated. In this letter, $\mu_{\texttt{OPT}}(\tilde{x}_t)$ is used as a bridge to connect $\mu_{\texttt{ALG},t}(\tilde{x}_t)$ and $\mu_{\texttt{OPT}}(x_t^\star)$. Some additional assumptions are needed.

Assumption 2 (Oracle Perturbation): The control law $\mu_{\texttt{OPT}}$ is Lipschitz continuous, i.e., there exists a constant $L_{\texttt{OPT}} > 0$ such that $\forall x', x'' \in \mathcal{X}$, it holds that

$$\|\mu_{\texttt{OPT}}\big(x'\big) - \mu_{\texttt{OPT}}\big(x''\big)\| \le L_{\texttt{OPT}}\|x' - x''\|. \quad (8)$$

Assumption 2 requires $\mu_{\texttt{OPT}}$ being at least continuous, which holds when f, ℓ, and F are all continuous, $\mathcal{U}$ is a polytope, and the solution $u_{0|t}^\star(x_t)$ is unique $\forall x_t \in \mathcal{X}_{\text{inv}}$ [1, Th. 2.7 and C.34]. Additional conditions to establish the Lipschitz continuity (8) require regularity of the optimization problem [5], [9]. Nonetheless, in the literature, the perturbation relation (8) is often either directly assumed [12], [18] or established with additional regularity properties [5], [8], [9].

Assumption 3 (Input Error Bound): There exist $\{\delta_t\}_{t=0}^{\infty}$, each of which is a mapping $\mathcal{X} \to \mathbb{R}_+$, such that

$$\|\mu_{\texttt{ALG},t}(x) - \mu_{\texttt{OPT}}(x)\| \le \delta_t(x), \ \forall x \in \mathcal{X}, \forall t \in \mathbb{N}, \quad (9)$$

where $\delta_t(x)$ is called the per-step control error.

Assumption 3 provides a quantitative bound for $\|\mu_{\texttt{ALG},t}(x) - \mu_{\texttt{OPT}}(x)\|$, and deriving an exact δ_t usually requires relaxations [5], [6], and thus (9) holds with inequality in general.

Remark 2 (Bounded Input Error): Assumption 3 is interpreted as the norm of the difference between the inputs generated using ALG and OPT is upper bounded by an explicit function of the state x, reflecting the input errors. One may also use a more conservative state-independent bound as $\delta_t^\star := \max_{x\in\mathcal{X}_{\text{inv}}} \delta_t(x)$. It should be noted that $\delta_t(x)$ (and thus also $\delta_t^\star$) is bounded, i.e., $\delta_t(x) \le \text{diam}(\mathcal{U})$. The specific form of $\delta_t(\cdot)$ depends on specific practical controllers.

According to (7) and (8), the regret upper bound depends on $\|x_t - x_t^\star\|$. In addition, it is obvious that the input error at any time t may also impact $\|x(k) - x^\star(k)\|$ for all $k \ge t+1$. Define the closed-loop disturbance $d_{f,t}(x)$ as

$$d_{f,t}(x) := f_{\texttt{ALG},t}(x) - f_{\texttt{OPT}}(x). \quad (10)$$

The closed-loop system in (5) can then be rewritten as

$$\tilde{x}_{t+1} = f_{\texttt{OPT}}\big(\tilde{x}_t\big) + d_{f,t}\big(\tilde{x}_t\big). \quad (11)$$

In addition, due to the Lipschitz continuity of f (cf. (2)) and Assumption 3, it holds that

$$\|d_{f,t}(x)\| \le L_{f,u}\|\mu_{\texttt{ALG},t}(x) - \mu_{\texttt{OPT}}(x)\| \le L_{f,u}\delta_t(x). \quad (12)$$

To study the state perturbation, Lipschitz continuity of (4) is helpful. Given that both the system and the oracle controller are Lipschitz continuous (see (2) and (8)), the closed-loop dynamics $f_{\texttt{OPT}}$ under the oracle controller $\mu_{\texttt{OPT}}$ is thus Lipschitz continuous on $\mathcal{X}_{\text{inv}}$, i.e., for all $x', x'' \in \mathcal{X}_{\text{inv}}$, there exists a constant $L_{f,\texttt{OPT}} \ge 0$ such that

$$\|f_{\texttt{OPT}}\big(x'\big) - f_{\texttt{OPT}}\big(x''\big)\| \le L_{f,\texttt{OPT}}\|x' - x''\|. \quad (13)$$

Based on the reformulated dynamics in (11) and the disturbance bound in (12), a quantitative bound on the state perturbation $\|\tilde{x}_t - x_t^\star\|$ can be established, which is given in the following lemma.

Lemma 1 (State Perturbation): For any time step $t \in \mathbb{N}$, the state perturbation $\|\tilde{x}_t - x_t^\star\|$ satisfies

$$\|\tilde{x}_t - x_t^\star\| \le L_{f,u} \sum_{i=0}^{t-1} L_{f,\texttt{OPT}}^{t-1-i}\delta_i(\tilde{x}_i), \quad (14)$$

where $L_{f,\texttt{OPT}}$ is given in (13), $L_{f,u}$ is given in (2), and $\delta_i(\cdot)$ is given in Assumption 3.

The proof of Lemma 1 is in Appendix A. Having established the state perturbation bound, the regret upper bound can be further streamlined.

Proposition 1 (Preparatory Regret Upper Bound): The regret Reg_T satisfies

$$\text{Reg}_T \le \sum_{t=0}^{T-1} \left[L_{l,u} + S_L \sum_{i=0}^{T-2-t} L_{f,\texttt{OPT}}^{i}\right]\delta_t\big(\tilde{x}_t\big), \quad (15)$$

where $S_L := (L_{\ell,x} + L_{\ell,u}L_{\texttt{OPT}})L_{f,u}$.

The proof of Proposition 1 is in Appendix C. To elaborate, (15) provides an upper bound of Reg_T in terms of the per-step control error $\delta_t(\tilde{x}_t)$, explicitly reflecting how suboptimal the practical controller ALG is. The error bound in (15) is also *consistent*, i.e., the bound degenerates to 0 when $\delta_t(\cdot) = 0$ for all $t \in \mathbb{N}$. We further impose an assumption about ALG, describing how ALG may improve online.

Assumption 4 (Algorithmic Improvement): There exists a *critical time step* $T_{\text{ct}} < +\infty$ such that

$$\delta_t(\tilde{x}_t) \leq \eta_{t-T_{\mathrm{ct}}}\left[\delta_{t-1}(\tilde{x}_{t-1}) + c_t(\Delta\tilde{x}_t)\right], \forall t \geq T_{\mathrm{ct}}, \tag{16}$$

where $\eta_{t-T_{\mathrm{ct}}} \in [0, 1]$, $c_t \in \mathcal{K}$, and $\Delta\tilde{x}_t = \|\tilde{x}_t - \tilde{x}_{t-1}\|$.

The inequality (16) captures a broad range of scenarios. First, $\eta_{t-T_{\mathrm{ct}}}$ can be interpreted as learning rates, characterizing how ALG improves and thus approximates OPT online. Besides, $c_t(\Delta\tilde{x}_t)$ specifies the additional cost due to state variation, and this type of cost typically exists in MPC with model learning [8] and real-time iteration [12], and in approximate MPC with interpolation errors [19]. It should be noted that (16) is valid after the critical time step, reflecting data collection procedures in, for example, online identification [17]. Based on Assumption 4, an upper bound of $\delta_t(\tilde{x}_t)$ can be provided, which is given in the following lemma:

Lemma 2: For $t \geq T_{\mathrm{ct}}$, $\delta_t(\tilde{x}_t)$ satisfies

$$\delta_t(\tilde{x}_t) \leq \left(\prod_{i=0}^{t-T_{\mathrm{ct}}} \eta_i\right)\mathrm{diam}(\mathcal{U}) + \sum_{i=0}^{t-T_{\mathrm{ct}}}\left(\prod_{j=i}^{t-T_{\mathrm{ct}}} \eta_j\right)\left[c_{T_{\mathrm{ct}}+i}\left(\Delta\tilde{x}_{T_{\mathrm{ct}}+i}\right)\right]. \tag{17}$$

Lemma 2 quantifies how the learning rates $\{\eta_i\}_{i=0}^{\infty}$ influence $\delta_t(\tilde{x}_t)$, and thus how $\{\eta_i\}_{i=0}^{\infty}$ determine the quality of ALG. Therefore, investigating the relationship between Reg_T and $\{\eta_i\}_{i=0}^{\infty}$ is necessary to quantify the suboptimality of ALG compared to OPT as well as to gain insights on designing ALG to achieve better control performance.

B. Lipschitz-Contractive Oracle Controller

We first consider the case where the closed-loop system (4) admits a contractive perturbation, i.e., $L_{f,\mathtt{OPT}} < 1$. In terms of analyzing the state perturbation, there is no significant difference between imposing the incremental stability (E-δ-ISS [12]) condition and the contractive perturbation condition.[2] In this case, (15) can be simplified as

$$\mathrm{Reg}_T \leq \underbrace{\left(L_{\ell,u} + \frac{S_L}{1 - L_{f,\mathtt{OPT}}}\right)}_{:=\bar{L}} \sum_{t=0}^{T-1} \delta_t(\tilde{x}_t). \tag{18}$$

As such, the regret scales linearly with $\sum_{t=0}^{T-1} \delta_t(\tilde{x}_t)$, and the next step is to evaluate the accumulated per-step control error along $\{\tilde{x}_t\}_{t=0}^{T-1}$. The main results discussing the requirement of the learning rates $\{\eta_i\}_{i=0}^{\infty}$ are given next.

Theorem 1 (Sublinear Regret): Assume that the closed-loop system (4) is Lipschitz-contractive, i.e., $L_{f,\mathtt{OPT}} < 1$.

1) If $c_t(\cdot) = 0$ and $\{\eta_i\}_{i=1}^{\infty}$ satisfy

$$\eta_i \leq \left(\frac{i}{i+1}\right)^{\alpha}, \text{ for } \alpha \in (0, 1), \tag{19}$$

then it holds that $\mathrm{Reg}_T = \mathcal{O}(T^{1-\alpha})$. In addition, if (19) holds for $\alpha = 1$, then $\mathrm{Reg}_T \leq \mathcal{O}(\log T)$.

2) If $c_t(\Delta x) \leq \bar{c}$ (i.e., bounded) and $\{\eta_i\}_{i=1}^{\infty}$ satisfy

$$\eta_i\left(1 + \frac{1}{\sum_{k=0}^{i-1}\prod_{j=k}^{i-1}\eta_j}\right) \leq \left(\frac{i}{i+1}\right)^{\alpha}, \quad \text{for } \alpha \in (0, 1), \tag{20}$$

then it holds that $\mathrm{Reg}_T = \mathcal{O}(T^{1-\alpha})$. In addition, if (20) holds for $\alpha = 1$, then $\mathrm{Reg}_T = \mathcal{O}(\log T)$.

The first case in Theorem 1 states that, when no state-variation cost exists, the regret can be sublinear if the learning rate satisfies a weak condition in the sense that the bound of the learning rate converges to 1 (i.e., no learning is required eventually). For example, with $\eta_0 = 1$ and $\eta_i = (i/i+1)^{\frac{1}{2}} (i = 1, 2, \ldots)$, the regret satisfies $\mathrm{Reg}_T = \mathcal{O}(\sqrt{T})$. Several existing results [2], [6], [8] fit our analysis framework, and the decreasing bound (19) is consistent with the general results in RLS identification, where the convergence rate is of $\mathcal{O}(T^{-\frac{1}{2}})$ [20], [21]. To showcase the regret behavior, a numerical example of a parametric linear system in [22] is simulated. The system model is given by

$$x_{t+1} = A(\theta)x_t + B(\theta)u_t + w_t,$$

where $A(\theta)$ and $B(\theta)$ are parameterized matrices and w_t are i.i.d. bounded disturbance. OPT uses the true θ^* whereas ALG applies an RLS estimator that generates estimates $\hat{\theta}_t$ online.[3] Specific code implementations can be accessed via https://github.com/lcrekko/mpc_reg_error_input. For this example, the closed-loop state is bounded, $\mu_{\mathtt{OPT}}$ is continuous and piecewise-affine, and the learning rate of $\delta_t(\tilde{x}_t)$ satisfies the time-varying decreasing behavior [2], [21] with $c_t(\cdot) = 0$ (see Remark 3). As such, Assumptions 1 to 4 are all solid. Fig. 1 presents the regret curve for 100 realizations of w_t, from which a sublinear behavior can be observed.

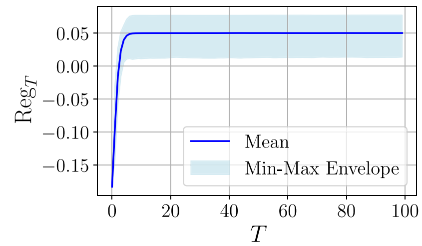

Fig. 1. Regret of MPC with online model learning.

On the other hand, for the more general case where the state-variation cost is bounded, a stronger requirement of the learning rate is needed to guarantee sublinear regret. A typical example is choosing $\eta_0 = 1$, $\eta_1 = \sqrt{2}/4$, and $\eta_i = (i\sqrt{i+1} - \sqrt{i(i+1)})(i^2-1)^{-1}$ (for $i = 2, 3, \ldots$), and then it holds that $\mathrm{Reg}_T = \mathcal{O}(\sqrt{T})$. Theorem 1 serves as a design guideline when a sublinear regret is desired such that, when input errors are present, the learning-based MPC controller is performing as well as its ideal counterpart in the long run. Specifically, one can directly aim to design a learning-based module to assist MPC such that (16) holds with the condition (19) or (20) satisfied. Moreover, the results in Theorem 1 can be used as sufficient conditions to check the validity of sublinear regret given a designed learning-based controller. Since a constant learning rate is commonly encountered and adopted in design [8], [12], it is also important to discuss the corresponding consequence in terms of the regret, which we present in the following theorem.

[2]The structure of the inequality in (14) is the same as the incremental input-to-state stability condition in [13, Definition 1].

[3]Unlike the robust MPC algorithm in [22], for simplicity, both OPT and ALG adopt certainty-equivalence principle and ignore w_t. The disturbance w_t is only added to excite the system for effective parameter identification [21].

Theorem 2 (Constant Learning Rates): Assume that the closed-loop system (4) is Lipschitz-contractive (i.e., $L_{f,\texttt{OPT}} < 1$) and $\eta_i = \eta^* \in (0,1)$, (for $i = 1, 2, \dots$).

1) If $c_t(\cdot) = 0$, then $\text{Reg}_T = \mathcal{O}(1)$.
2) If $c_t(\Delta x) \le \bar{c}$ (i.e., bounded), then $\text{Reg}_T = \mathcal{O}(T)$.

Essentially, Theorem 2 highlights that i) the regret is *finite* when c_t is not present and ii) the regret is *linear* when c_t is bounded. In other words, when c_t is non-zero, a constant learning rate is insufficient to ensure sublinear regret, thereby necessitating the more stringent condition (20). As this condition requires the learning rate to asymptotically approach zero–an inherently more restrictive requirement–we present potential refinements in the subsequent remarks.

Remark 3 (Uniform Error Bound): Following Remark 2, the worst-case perturbation $\delta^\star$ can quantify the quality of $\texttt{ALG}$. Accordingly, (16) in Assumption 4 degenerates to

$$\delta_t^\star \le \eta_{t-T_{\text{ct}}}\delta_{t-1}^\star, \forall t \ge T_{\text{ct}}, \tag{21}$$

and the result from the first case of Theorem 1 is thus directly applicable. If (21) can be guaranteed with $\eta_i \in (0,1)$, less stringent requirements on the learning rates (19) can achieve sublinear regret. For example, if the model (1) is parametric, i.e., $x_{k+1} = f(x_k, u_k; \theta)$ with $\theta \in \Theta \subset \mathbb{R}^p$ and the true parameter being θ^*, then the parametric control law is $\mu(x;\theta)$. It has been established that $\|\mu(x;\theta') - \mu(x;\theta'')\| \le L\|\theta' - \theta''\|$ for all $\theta', \theta'' \in \Theta$ and some L under mild conditions[4] [5], [9]. Therefore, consider $\mu_{\texttt{OPT}}(x) = \mu(x;\theta^*)$ and $\mu_{\texttt{ALG},t}(x) = \mu(x;\hat{\theta}_t)$ with $\hat{\theta}_t$ being the online estimated parameter, a state-independent bound $\delta_t^\star = L\|\theta^* - \hat{\theta}_t\|$ is thus valid.

Remark 4 (Refinements Through Stability): In cases where $\texttt{ALG}$ stabilizes the system at an equilibrium x_{eq}, then the cost due to the state variation eventually vanishes, leading to weaker learning rate requirements (19). Nonetheless, achieving stability using $\texttt{ALG}$ can be challenging [2].

C. Lipschitz-Noncontractive Oracle Controller

Next, we consider the case where the closed-loop system in (4) satisfies a non-contractive perturbation, i.e., $L_{f,\texttt{OPT}} \ge 1$. This scenario is also frequently encountered in practice, and a naive example is linear systems $x_{t+1}^\star = f_{\texttt{OPT}}(x_t^\star) = Ax_t^\star$ with $\|A\| \ge 1$. Given $L_{f,\texttt{OPT}} = 1$, it holds that

$$\text{Reg}_T \le \sum_{t=0}^{T-1}\left[L_{\ell,u} + S_L(T-1-t)\right]\delta_t(\tilde{x}_t), \tag{22}$$

and when $L_{f,\texttt{OPT}} > 1$, we again adopt (15) to derive the results. Since $\delta_0(\tilde{x}_0)$ is non-zero in general, it can only be shown that $\text{Reg}_T = \mathcal{O}(T)$ when $L_{f,\texttt{OPT}} = 1$ and $\text{Reg}_T = \mathcal{O}\left(L_{f,\texttt{OPT}}^T\right)$ when $L_{f,\texttt{OPT}} > 1$. Thus, even though the true regret can be sublinear, no sublinear regret can be concluded theoretically when the oracle controller cannot yield a contractive closed-loop system. This limitation explains why a stabilizable linear dominant system is considered in [8] and why the incremental stability is assumed in [12]. Leveraging the results in Section IV-B, it is demonstrated that contraction is crucial to obtain meaningful suboptimality guarantees. Similar insights have been reported in other works using other contraction notions, where stability and safety are the primary concerns [18].

[4]See the assumptions on linear independence constraint qualification and strong second-order sufficient condition in [5], [9].

V. Conclusions & Future Work

This letter has presented a high-level regret analysis of MPC for Lipschitz continuous nonlinear systems with imperfect inputs, revealing the suboptimality of a broad range of MPC strategies with imperfect inputs. Specifically, considering the oracle controller being Lipschitz-contractive, sufficient conditions on the regulated learning rates of the practical controller are provided such that the dynamic regret is sublinear. Besides, given a constant learning rate, it is shown that sublinear regret cannot be achieved within the used analysis framework if the state-variation cost persists. Finally, it is demonstrated that, in our analysis framework, no sublinear regret can be derived when the oracle controller is non-contractive.

Future working directions include i) designing specific MPC controllers that meet the learning-rate conditions, ii) investigating the performance of learning-based output-tracking MPC, and iii) extending the framework to general optimal control strategies and problems with explicit state constraints.

Appendix

A. Proof of Lemma 1

Proof: The proof is based on induction. Let $t = 0$, $\|x_0 - x_0^\star\| = \|x - x\| = 0$ as the sum $\sum_0^{-1}$ is 0 by definition, (14) is valid. Assume that (14) holds for $t = k$, i.e.,

$$\|\tilde{x}_k - x_k^\star\| \le L_{f,u}\sum_{i=0}^{k-1} L_{f,\texttt{OPT}}^{k-1-i}\delta_i(\tilde{x}_i). \tag{23}$$

For $t = k+1$, it holds that

$$\begin{aligned}\|x_{k+1} - x_{k+1}^\star\| &\le \|f_{\texttt{OPT}}(\tilde{x}_k) - f_{\texttt{OPT}}\left(x_k^\star\right)\| + \|d_{f,k}(\tilde{x}_k)\| \\ &\overset{(12),(13)}{\le} L_{f,\texttt{OPT}}\|\tilde{x}_k - x_k^\star\| + L_{f,u}\delta_t(\tilde{x}_k)^{(23)} \\ &\le L_{f,u}\sum_{i=0}^{k} L_{f,\texttt{OPT}}^{k-i}\delta_i(\tilde{x}_i).\end{aligned}$$

Thus, the proof is completed by induction. ■

B. Proof of Lemma 2

Proof: The proof is similar to the proof of Lemma 2 by induction, and the details are omitted for brevity. The final bound employs $\delta_{T_{\text{ct}}-1}(\tilde{x}_{T_{\text{ct}}-1}) \le \text{diam}(\mathcal{U})$. ■

C. Proof of Proposition 1

Proof: First, the mismatched input error $\mu_{\texttt{ALG},t}(\tilde{x}_t) - \mu_{\texttt{OPT}}(x_t^\star)$ can be bounded as

$$\begin{aligned}&\|\mu_{\texttt{ALG},t}(\tilde{x}_t) - \mu_{\texttt{OPT}}\left(x_t^\star\right)\| \\ &\le \|\mu_{\texttt{ALG},t}(\tilde{x}_t) - \mu_{\texttt{OPT}}(\tilde{x}_t)\| + \|\mu_{\texttt{OPT}}(\tilde{x}_t) - \mu_{\texttt{OPT}}\left(x_t^\star\right)\| \\ &\overset{(8),(9)}{\le} \delta_t\left(\tilde{x}_t\right) + L_{\texttt{OPT}}\|\tilde{x}_t - x_t^\star\|.\end{aligned} \tag{24}$$

Substituting (24) into (7) yields

$$\text{Reg}_T(x) \le L_{\ell,u}\sum_{t=0}^{T-1}\delta_t\left(\tilde{x}_t\right) + \left(L_{\ell,x} + L_{\ell,u}L_{\texttt{OPT}}\right)\sum_{t=0}^{T-1}\|\tilde{x}_t - x_t^\star\|. \tag{25}$$

Finally, we obtain (15) by substituting (14) into (25) and rearranging the terms property. ■

D. Proof of Theorem 1

Proof: The proof starts with part 1). From (17), assigning $c_{T_{\rm ct}+i}(\cdot) = 0$ yields

$$\delta_t\big(x(t)\big) \leq \left(\prod_{i=0}^{t-T_{\rm ct}} \eta_i\right) \mathrm{diam}(\mathcal{U}). \tag{26}$$

Based on (18), substituting (26) leads to

$$\mathrm{Reg}_T \overset{(19)}{\leq} \mathrm{diam}(\mathcal{U})\bar{L}\left[T_{\rm ct} + \eta_0 \sum_{t=1}^{T-T_{\rm ct}} \left(\frac{1}{t}\right)^{\alpha}\right]. \tag{27}$$

For $\alpha \in (0, 1)$, (27) implies that $\mathrm{Reg}_T(x) = \mathcal{O}\big(T^{1-\alpha}\big)$. On the other hand, for $\alpha = 1$, (27) entails that $\mathrm{Reg}_T(x) = \mathcal{O}(\log T)$.

In terms of part 2), given that both $c_t(\cdot)$ and ω_t are bounded, by substituting (17) into (18), the following result can be obtained:

$$\mathrm{Reg}_T \leq \mathrm{diam}(\mathcal{U})\bar{L}\Bigg[T_{\rm ct} + \sum_{t=T_{\rm ct}}^{T}\left(\prod_{i=0}^{t-T_{\rm ct}} \eta_i\right) + \sum_{t=T_{\rm ct}}^{T}\sum_{i=0}^{t-T_{\rm ct}}\left(\prod_{j=i}^{t-T_{\rm ct}} \eta_j\right)\Bigg]. \tag{28}$$

Since $\eta_i > 0$, it is obvious that $1 + (\sum_{k=0}^{i-1}\prod_{j=k}^{i-1}\eta_j)^{-1} > 1$, indicating that (20) implies (19). Thus, following (27), we again have

$$\sum_{t=T_{\rm ct}}^{T}\left(\prod_{i=0}^{t-T_{\rm ct}} \eta_i\right) \leq \eta_0 \sum_{t=1}^{T-T_{\rm ct}}\left(\frac{1}{t}\right)^{\alpha}. \tag{29}$$

Moreover, from (20), it holds that

$$\sum_{t=T_{\rm ct}}^{T}\sum_{i=0}^{t-T_{\rm ct}}\left(\prod_{j=i}^{t-T_{\rm ct}} \eta_j\right) \leq \eta_0 \sum_{t=1}^{T-T_{\rm ct}}\left(\frac{1}{t}\right)^{\alpha}. \tag{30}$$

Substituting (29) and (30) into (28) yields

$$\mathrm{Reg}_T(x) \leq \mathrm{diam}(\mathcal{U})\bar{L}\Bigg[T_{\rm ct} + \eta_0\left(1 + \frac{\bar{c}}{\mathrm{diam}(\mathcal{U})}\right)\sum_{t=1}^{T-T_{\rm ct}}\left(\frac{1}{t}\right)^{\alpha}\Bigg]. \tag{31}$$

Similarly, given (31), it holds that $\mathrm{Reg}_T(x) = \mathcal{O}\big(T^{1-\alpha}\big)$ for $\alpha \in (0, 1)$ and $\mathrm{Reg}_T(x) \leq \mathcal{O}(\log T)$ for $\alpha = 1$. ■

E. Proof of Theorem 2

Proof: When $\eta_i = \eta^*$, it is easy to verify that

$$\sum_{t=T_{\rm ct}}^{T}\left(\prod_{i=0}^{t-T_{\rm ct}} \eta_i\right) \leq \frac{\eta^*}{1-\eta^*} = \mathcal{O}(1), \tag{32a}$$

$$\sum_{t=T_{\rm ct}}^{T}\sum_{i=0}^{t-T_{\rm ct}}\left(\prod_{j=i}^{t-T_{\rm ct}} \eta_j\right) \leq \frac{T\eta^*}{1-\eta^*} = \mathcal{O}(T). \tag{32b}$$

The rest of the proof is similar to that of Theorem 1 in Appendix D. ■

References

[1] J. Rawlings, D. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. San Francisco, CA, USA: Nob Hill Publ., 2017.

[2] H. Zhou and V. Tzoumas, "Simultaneous system identification and model predictive control with no dynamic regret," 2024, *arXiv:2407.04143*.

[3] J. Maciejowski and C. Jones, "MPC fault-tolerant flight control case study: Flight 1862," *IFAC Proc. Vol.*, vol. 36, no. 5, pp. 119–124, 2003.

[4] C. Liu, S. Shi, and B. De Schutter, "Stability and performance analysis of model predictive control of uncertain linear systems," in *Proc. IEEE 63rd Conf. Decision Control (CDC)*, 2024, pp. 7356–7362.

[5] C. Liu, S. Shi, and B. De Schutter, "Certainty-equivalence model predictive control: Stability, performance, and beyond," 2024, *arXiv:2412.10625*.

[6] A. Karapetyan, E. Balta, A. Iannelli, and J. Lygeros, "On the finite-time behavior of suboptimal linear model predictive control," in *Proc. 62nd IEEE Conf. Decision Control (CDC)*, 2023, pp. 5053–5058.

[7] M. Lorenzen, F. Allgöwer, and M. Cannon, "Adaptive model predictive control with robust constraint satisfaction," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 3313–3318, 2017.

[8] I. Dogan, M. Shen, and A. Aswani, "Regret analysis of learning-based MPC with partially-unknown cost function," *IEEE Trans. Autom. Control*, vol. 69, no. 5, pp. 3246–3253, May 2024.

[9] Y. Lin, T. Hu, G. Qu, T. Li, and A. Wierman, "Bounded-regret MPC via perturbation analysis: Prediction error, constraints, and non-linearity," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 35, 2022, pp. 36174–36187.

[10] S. H. Son, J. W. Kim, T. H. Oh, D. H. Jeong, and J. M. Lee, "Learning of model-plant mismatch map via neural network modeling and its application to offset-free model predictive control," *J. Process Control*, vol. 115, pp. 112–122, Jul. 2022.

[11] I. Schimperna and L. Magni, "Robust offset-free constrained model predictive control with long short-term memory networks," *IEEE Trans. Autom. Control*, vol. 69, no. 12, pp. 8172–8187, Dec. 2024.

[12] A. Karapetyan, E. Balta, A. Iannelli, and J. Lygeros, "Closed-loop finite-time analysis of suboptimal online control," *IEEE Trans. Autom. Control*, early access, Feb. 7, 2025, doi: 10.1109/TAC.2025.3539988.

[13] D. N. Tran, B. S. Rüffer, and C. M. Kellett, "Convergence properties for discrete-time nonlinear systems," *IEEE Trans. Autom. Control*, vol. 64, no. 8, pp. 3415–3422, Aug. 2019.

[14] Y. Lin, G. Goel, and A. Wierman, "Online optimization with predictions and non-convex losses," *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 4, no. 1, pp. 1–32, 2020.

[15] J. Rawlings, D. Angeli, and C. Bates, "Fundamentals of economic model predictive control," in *Proc. IEEE 51st Conf. Decision Control (CDC)*, 2012, pp. 3851–3861.

[16] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*. Cambridge, U.K.: Cambridge Univ., 2017.

[17] D. Muthirayan, J. Yuan, D. Kalathil, and P. Khargonekar, "Online learning for predictive control with provable regret guarantees," in *Proc. IEEE 61st Conf. Decision Control (CDC)*, 2022, pp. 6666–6671.

[18] C. Dawson, S. Gao, and C. Fan, "Safe control with learned certificates: A survey of neural Lyapunov, barrier, and contraction methods for robotics and control," *IEEE Trans. Robot.*, vol. 39, no. 3, pp. 1749–1767, Jun. 2023.

[19] A. Rose, P. Schaub, and R. Findeisen, "Safe and high-performance learning of model predictive control using kernel-based interpolation," 2024, *arXiv:2410.06771*.

[20] B. Polyak and A. Juditsky, "Acceleration of stochastic approximation by averaging," *SIAM J. Control Optim.*, vol. 30, no. 4, pp. 838–855.

[21] L. Ljung and T. Söderström, *Theory and Practice of Recursive Identification*. Cambridge, MA, USA: MIT press, 1983.

[22] M. Lorenzen, M. Cannon, and F. Allgöwer, "Robust MPC with recursive model update," *Automatica*, vol. 103, pp. 461–471, May 2019.