A Bilevel Approach for Nonlinear Optimal Control Problems with Free Final Time

Alberto De Marchi[a,*], Matthias Gerdts[a]

[a] *Bundeswehr University Munich*
Werner-Heisenberg-Weg 39, 85577 Neubiberg, Germany

Abstract

This paper discusses an interpretation of variable time-horizon nonlinear optimal control problems as bilevel optimization problems and addresses a method for their approximate solution. The developed method formulates a bilevel optimization problem seeking, at the upper level, the optimal final time and optimal control and corresponding state at the lower level. The latter can be solved by classic tools for nonlinear optimal control problems. Instead, necessary and sufficient optimality conditions for the free final time are given, based on the associated Hamilton function. Then, any root-finding strategy can be exploited to tackle the upper-level problem through the corresponding necessary conditions. Furthermore, for a specific method found in the literature, it was possible to seamlessly couple the two optimization levels and to use a second-order derivative-based solver for the upper-level problem. Numerical results are obtained for a problem from vehicle trajectory planning that involves a trade-off between control effort and final time. The results show good convergence behaviour even for inexact lower-level solutions.

Keywords: Optimal control, bilevel optimization, free final time, nonlinear systems, Hamilton function
2010 MSC: 49J15, 49M05, 90C52

1. INTRODUCTION

The field of optimal control is becoming more and widely common in automatic control applications. This paper addresses variable finite time-horizon nonlinear optimal control problems (OCP) associated with continuous-time, deterministic, autonomous, nonlinear controlled dynamical systems, coupled boundary conditions, autonomous nonlinear Bolza-type cost functionals and variable finite time-horizon. This class of problems has been widely studied and discussed in the literature, see e.g. [1, 2, 3]. The contribution of this work is the interpretation of variable time-horizon OCPs as bilevel optimization problems (BOP) and the resulting solution approach. This allows to avoid any time-transformation technique, e.g. [3, 4, 5], while rather to exploit the problem's structure. A similar two-level procedure has been developed in [6] for time optimal MPC applications; therein, however, time optimality is strictly prioritized over other objectives. Usually, solution approaches for BOPs aim at reducing the bilevel structure into a possibly equivalent single stage optimization problem, e.g. [7]. Herein, we reverse the perspective and formulate a single-objective BOP equivalent to the original OCP. On the traces of Romans' motto *divide et impera*, both the optimization levels are seemingly easier to solve than the initial problem. In fact, the so-called upper-level problem seeks the optimal time-horizon for the control task, and thus it has just a scalar decision variable, while the fixed finite time-horizon optimal control for the given nonlinear dynamics and boundary conditions is managed at the lower level. To the best of authors' knowledge, this bilevel viewpoint has been introduced for linear systems only, in [8], and hence it is a novel contribution in the field of nonlinear optimal control. This work aims at introducing this perspective, discussing advantages and limitations and reporting results on a numerical example; nonetheless, this is far from being a complete analysis and critique of the method. Similarly to previous work, necessary and sufficient conditions for the upper-level optimality are founded on the condition of vanishing Hamilton function along solutions to variable finite time-horizon OCPs [1, 3]. Instead, any standard numerical method for fixed finite time-horizon OCPs can be adopted for tackling the lower-level optimization problem. In this work, for the numerical example, we focus on a method that makes relatively easy to couple the upper and the lower level. As pointed out in [8], the bilevel approach helps improving robustness and stability of the solution process, and provides an intuitive interpretation in the context of model predictive control (MPC), namely a tracking procedure for the optimal final time. From the theoretical viewpoint, the suggested perspective exploits a property of the Hamilton function, which is usually not easy, if possible, with di-

*Corresponding author
Email addresses: `alberto.demarchi@unibw.de` (Alberto De Marchi), `matthias.gerdts@unibw.de` (Matthias Gerdts)
[0]Declarations of interest: none.

rect methods after time- or spatial transformation. Also, one could consider a very general setting, namely all the fixed finite time-horizon OCPs that can be solved with standard numerical methods [9, 3]. Focusing this work on an extension toward variable finite time-horizon problems, the class of problems is restricted (nonlinear autonomous dynamics and linear-quadratic autonomous cost) for the sake of simplicity. However, we highlight that for nonautonomous OCPs, i.e. those with either nonautonomous dynamics or time-dependent cost functional, just a detail changes, as discussed in Section 3 about Eq. 5.

The paper is organized as follows. Section 2 states the problem with some standing assumptions, delineates the solution approach and emphasizes the bilevel perspective. Section 3 analyses the free final time problem from the upper-level viewpoint and examines how the bilevel problem can be iteratively solved. In Section 4 a specific method, namely the MASRE method [9, 10], is discussed as a solver for the lower-level problem and the coupling with the upper-level problem is analysed. Section 5 validates the proposed approach numerically on a trajectory planning problem, showing effectiveness and limitations of the proposed algorithm. Section 6 concludes the paper and presents ideas for future research.

2. Problem formulation

Consider a time interval $[0, T]$, with final time $T > 0$, a nonlinear autonomous control-nonaffine state differential equation

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t)) \tag{1}$$

for a.e. $t \in [0, T]$, boundary conditions

$$\mathbf{b}\left(\mathbf{x}(0), \mathbf{x}(T)\right) = \mathbf{0} \tag{2}$$

and a cost functional J, defined by

$$J(\mathbf{x}, \mathbf{u}, T) := \int_0^T \left[\omega + \ell(\mathbf{x}(\tau), \mathbf{u}(\tau))\right] \, \mathrm{d}\tau \tag{3}$$

Herein, for any time t, $\mathbf{x}(t) \in \mathbb{R}^{n_x}$ denotes the state and $\mathbf{u}(t) \in \mathbb{R}^{n_u}$ the control input. Also, let functions $\mathbf{f} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x}$, $\mathbf{b} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_x} \to \mathbb{R}^{n_b}$, $\ell : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}_{\geq 0}$ and parameter $\omega > 0$ be given. Let us define $\Gamma = (0, +\infty)$ and, given a final time $T \in \Gamma$, $X(T) := W^{1,\infty}([0, T], \mathbb{R}^{n_x})$ the Sobolev space of vector-valued absolutely continuous functions with essentially bounded first derivative on $[0, T]$, and $U(T) := L^\infty([0, T], \mathbb{R}^{n_u})$ the vector space of vector-valued essentially bounded measurable functions on $[0, T]$. The original, nonlinear autonomous free end time OCP, or OOCP, dealt with in this paper is the following:

Problem 1 (OOCP)**.** *Find a final time $T \in \Gamma$, a control $\mathbf{u} \in U(T)$, and a state $\mathbf{x} \in X(T)$, minimizing the cost functional J in (3) while satisfying the constraints (1)–(2).*

The bilevel approach stems from decoupling the free final time optimization, on one side, and the control input optimization, with the associated state. To this end, let us cast the OOCP into an equivalent BOP, whose lower-level problem (LLP), given a final time $T \in \Gamma$, reads:

Problem 2 (LLP$_T$)**.** *Find a control $\mathbf{u} \in U(T)$, and a state $\mathbf{x} \in X(T)$, minimizing the cost functional J in (3) while satisfying the constraints (1)–(2).*

The upper-level problem (ULP) reads:

Problem 3 (ULP)**.** *Find a final time $T \in \Gamma$ minimizing the cost function $\tilde{J}(T) := J(\mathbf{x}_T, \mathbf{u}_T, T)$, with $(\mathbf{x}_T, \mathbf{u}_T)$ solving the LLP$_T$.*

In a sense, the LLP is to be solved just to evaluate the upper-level cost function $\tilde{J}$, that is the cost functional as seen from the ULP. Note that the ULP has only a scalar optimization variable, while the LLP is a standard fixed finite time-horizon OCP. For the ULP, instead, we consider necessary and sufficient optimality conditions and derive a simple yet seemingly effective numerical method for its solution. The approach goes back to that of [8] for linear time-invariant system and extends it to deal with nonlinear dynamics. We remark that the single-objective BOP consisting of Problems 2 and 3 is equivalent to the OOCP.

In this paper the following standing assumptions are considered:

Assumption 1. *For any given final time $T \in \Gamma$, the LLP$_T$ admits a solution;*

Assumption 2. *The set of optimal final times for the OOCP has zero measure;*

Hence, existence of a solution to the OOCP comes from Assumption 1 and positiveness of ω. Equivalently, in the case ω is not strictly positive, the set Γ must be a closed interval [8].

3. Optimality conditions for the free final time

This section addresses necessary and sufficient optimality conditions for the free final time in the OOCP, stemming from the vanishing condition on the Hamilton function for free time OCPs. Such approach stretches previous work on linear time-invariant systems and considers nonlinear autonomous systems.

3.1. Necessary and sufficient (local) optimality conditions

Let us denote $(\mathbf{x}^\star, \mathbf{u}^\star, T^\star, \boldsymbol{\lambda}^\star, \boldsymbol{\eta}^\star)$ a solution to the OOCP with the corresponding Lagrange multipliers, and $\mathcal{H}$ the Hamilton function associated with the OOCP.

$$\mathcal{H}(\mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}) := \omega + \ell(\mathbf{x}, \mathbf{u}) + \boldsymbol{\lambda}^\top \mathbf{f}(\mathbf{x}, \mathbf{u}) \tag{4}$$

Also, let $(\mathbf{x}^\star_T, \mathbf{u}^\star_T, \boldsymbol{\lambda}^\star_T, \boldsymbol{\eta}^\star_T)$ denote a solution to the LLP$_T$, for any given $T \in \Gamma$. The idea behind the bilevel approach and the method developed to tackle the ULP stems from the following observations:

(i) For an autonomous OCP with free final time, the Hamilton function vanishes almost everywhere in time along a solution [3, Thm. 7.1.8]; namely, it holds $\mathcal{H}(\mathbf{x}^\star(t), \mathbf{u}^\star(t), \boldsymbol{\lambda}^\star(t)) = 0$ for a.e. $t \in [0, T^\star]$.

(ii) There exists a solution, say $(\mathbf{x}^\star_{T^\star}, \mathbf{u}^\star_{T^\star}, \boldsymbol{\lambda}^\star_{T^\star}, \boldsymbol{\eta}^\star_{T^\star})$, to the LLP$_{T^\star}$ such that $\mathbf{x}^\star = \mathbf{x}^\star_{T^\star}$, $\mathbf{u}^\star = \mathbf{u}^\star_{T^\star}$, $\boldsymbol{\lambda}^\star = \boldsymbol{\lambda}^\star_{T^\star}$ and $\boldsymbol{\eta}^\star = \boldsymbol{\eta}^\star_{T^\star}$.

(iii) For any $T \in \Gamma$, the Hamilton function $\mathcal{H}$ is also associated with the LLP$_T$.

Inspired by these observations, we define $h : \Gamma \to \mathbb{R}$ as

$$h(T) := \mathcal{H}\left(\mathbf{x}^\star_T(\tau), \mathbf{u}^\star_T(\tau), \boldsymbol{\lambda}^\star_T(\tau)\right) \tag{5}$$

for any arbitrary $\tau \in [0, T]$. In principle, the value of τ can be freely chosen because for autonomous OCPs the Hamilton function is constant along a solution [1, 3]; instead, for nonautonomous problems, the vanishing condition holds only at the final time [3]; hence, one must consider $\tau = T$. Considering observations (i)–(iii) and definition (5), the necessary optimality condition

$$h(T^\star) = 0 \tag{6}$$

follows. Thus, exploiting this condition, we may replace the ULP with the surrogate problem of seeking a $T \in \Gamma$ such that $h(T) = 0$. In fact, the vanishing condition is only necessary in that it is also satisfied at local minima and maxima; however, one can a posteriori check for optimality, as we shall detail next. Also, we stress that it is common practice to exploit necessary conditions in numerical methods for optimization, even though these alone cannot guarantee optimality [2, 3, 11]. Nonetheless, depending on the problem, it might be possible to tighten some of the aforementioned conditions, obtaining stronger results.

The necessary vanishing condition on the Hamilton function is tightly connected with the necessary stationarity condition on the reduced cost function $\tilde{J}$ introduced in Problem 3. In fact, with Assumption 1, the identity $\tilde{J}'(T) = h(T)$ for any $T \in \Gamma$, obtained in [8, Thm. 10], can be easily extended to fit nonlinear systems. This relationship gives an interpretation of the vanishing condition on the Hamiltonian. Moreover, based on that, a sufficient condition for optimality can be derived: a final time $T^\star$, such that $h(T^\star) = 0$, is (locally) optimal, that is, (locally) minimizes the (reduced) cost function, if $h'(T^\star) > 0$.

Let us shift the focus on the ULP-LLP coupling. Many different methods exist to deal with the (fixed finite time-horizon) LLP, e.g. direct methods with full discretization or collocation, indirect methods and dynamic programming [1, 2, 3, 11]. Let us consider a generic method for the LLP, say a lower-level method (LLM), and, for any given $T \in \Gamma$, assume it is able to find a solution to the LLP$_T$, possibly through a sequence of sub-problems, say LLP$^{[1]}_T$, LLP$^{[2]}_T$, ..., eventually converging to the original LLP$_T$. In practice, this class of methods includes those exploiting relaxation, penalization, continuation and iterative approximation techniques; indeed, direct methods can be considered as one-iteration methods. In fact, the MASRE method discussed in Section 4 belongs to this class too. We denote LLP$^\infty_T$ the problem identical to the actual LLP$_T$ (possibly up to a certain numerical tolerance); the number of sub-problems might be 1 as well as infinite, depending on the adopted method, prescribed tolerances, numerical precision and sought exactness.

For a given $T \in \Gamma$, we denote $(\mathbf{x}^\star_T, \mathbf{u}^\star_T, \boldsymbol{\lambda}^\star_T, \boldsymbol{\eta}^\star_T)$ a solution to the LLP$_T$, and $\left(\mathbf{x}^{[k]}_T, \mathbf{u}^{[k]}_T, \boldsymbol{\lambda}^{[k]}_T, \boldsymbol{\eta}^{[k]}_T\right)$ a solution to the LLP$^{[k]}_T$, namely the k-th sub-problem. Let $\mathcal{H}^{[k]}_{\mathrm{LL}}$ be the Hamilton function associated with the LLP$^{[k]}_T$, possibly dependent on time. Further useful observations concerning the coupling are in order:

(iv) if the LLM (globally) converges to a solution of the LLP$_T$, then there exist solutions $(\mathbf{x}^\infty_T, \mathbf{u}^\infty_T, \boldsymbol{\lambda}^\infty_T, \boldsymbol{\eta}^\infty_T)$ and $(\mathbf{x}^\star_T, \mathbf{u}^\star_T, \boldsymbol{\lambda}^\star_T, \boldsymbol{\eta}^\star_T)$ such that $\mathbf{x}^\infty_T = \mathbf{x}^\star_T$, $\mathbf{u}^\infty_T = \mathbf{u}^\star_T$, $\boldsymbol{\lambda}^\infty_T = \boldsymbol{\lambda}^\star_T$ and $\boldsymbol{\eta}^\infty_T = \boldsymbol{\eta}^\star_T$;

(v) for any given $T \in \Gamma$, it holds

$$\mathcal{H}\left(\mathbf{x}^\star_T(t), \mathbf{u}^\star_T(t), \boldsymbol{\lambda}^\star_T(t)\right) = \mathcal{H}^\infty_{\mathrm{LL}}\left(t, \mathbf{x}^\infty_T(t), \mathbf{u}^\infty_T(t), \boldsymbol{\lambda}^\infty_T(t)\right) \tag{7}$$

for every $t \in [0, T]$.

Based on these observations, and analogously to (5), we define $h_{\mathrm{LL}} : \Gamma \to \mathbb{R}$ as

$$h_{\mathrm{LL}}(T) := \mathcal{H}^\infty_{\mathrm{LL}}(\tau, \mathbf{x}^\infty_T(\tau), \mathbf{u}^\infty_T(\tau), \boldsymbol{\lambda}^\infty_T(\tau)) \tag{8}$$

for any $\tau \in [0, T]$. Then, from conditions in (6) and (7) and definition (8), we have the necessary condition

$$h_{\mathrm{LL}}(T^\star) = 0 \tag{9}$$

that can be adopted as a surrogate problem. This makes it possible in principle to take advantage of information generated at, and provided by, the LLP to solve the ULP.

3.2. Scalar root-finding methods

This section aims at exploiting the necessary conditions (6), (9), and hence addresses root-finding algorithms for solving scalar real nonlinear equations.

Assumption 3. *Function $h : \Gamma \to \mathbb{R}$ is continuous in a neighborhood of every optimal final time $T^\star$.*

The choice of the upper-level method (ULM) might depend on its own properties and on the particular application, but it is independent on the LLM. In fact, the evaluation of function h and, if possible, of its derivatives, might depend on the LLM. Depending on the availability and the order of the derivatives of h, high-order methods can be adopted as ULM, yielding high convergence rates, e.g. Newton's and Halley's method [12, 13]; otherwise, bracketing and derivative-free methods, e.g. bisection, secant and Brent's method [14], come into play. Usually, these algorithms require an initial guess (sufficiently close to a root) or an interval (the function attaining different signs at the extrema), in order to start the iterative refining process. Notice that Assumption 2 is introduced for the purpose of facilitating the root-finding procedure and obtaining a more robust solution [8, Remark 5].

Remark 1. *Throughout this paper the free final time is considered unconstrained, even though it can attain finite positive values only. One could extend the present method to the case of box-constrained final time, namely having* $\Gamma = [T_l, T_u]$, *with* $0 < T_l < T_u < +\infty$. *It has been shown that a projected step is equivalent to accounting for the additional Lagrange multipliers and the associated complementarity conditions [8, Appendix B]. In practice, once the refined estimate is computed (by any method), it is projected onto the interval of feasible final times.*

Remark 2. *Within this work the MASRE method [10] has been adopted, in order to exploit the smoothness of the underlying linear-quadratic approximations, Section 4. Doing so, one can adopt gradient-based root-finding methods. Results obtained with this approach are reported in the numerical tests, Section 5. However, the adoption of derivative-free algorithms at the upper-level, possibly with other LLMs, e.g. OCPID-DAE1 [15] and CasADi [16], is subject of ongoing research.*

3.3. Algorithm

In Algorithm 1 the bilevel solver for free time OCPs is outlined and its key components highlighted. Therein, $\mathbf{S}$ denotes a solution to a LLP, collecting state, control, multipliers and possibly other variables, depending on the LLM. Similarly, $\mathbf{H}$ contains information about the Hamiltonian, that is, its value and possibly its derivatives with respect to the final time. There are three main subroutines:

- LLMStep returns a solution to a LLP using the LLM, given the problem instance, a final time and an initial guess, possibly from a previous solution;
- HamiFun evaluates the Hamilton function (8) and possibly its derivatives (say, for $\tau = 0$) at the current final time, given the problem instance and a lower-level solution;
- ULMStep generates a refined estimate of the optimal final time, performing a possibly projected step, based on the ULM, given the actual final time, the available information about the Hamiltonian and the feasible set; see Remark 1.

The structure of Algorithm 1 resembles the proposed bilevel perspective, and the main steps refer to the lower and upper level and their coupling. The inner loop accounts for the LLP through the LLM, while the ULP is tackled via the ULM in the outer loop, with the Hamiltonian as an interface across the levels.

Suitable termination criteria should be formulated for both, the inner and the outer loop. These may comprise a maximal number of iterations (N_{UL}, N_{LL}) and a prescribed minimum tolerance (ϵ_{UL}, ϵ_{LL}) on successive iterations, in terms of relative and absolute difference. For the ULP, both the final time and the value of the Hamilton function could be considered. For the LLP, instead, these criteria may depend on the LLM; one could compare successive solutions, states or controls, possibly at specific time points [9, 10]. For automatic control applications, e.g. MPC, one could also stop the outer loop once the initial control input has converged.

Algorithm 1 Bilevel solver for free final time OCPs

```
Input: Γ, T^[0], S^[0,0], ε_UL, ε_LL, N_UL, N_LL
Output: T*, S*, H
  for j = 1, ..., N_UL do                                  // upper level
      for k = 1, ..., N_LL do                              // lower level
          S^[j-1,k] ← LLMStep(T^[j-1], S^[j-1,k-1])
          if ||S^[j-1,k] − S^[j-1,k-1]|| < ε_LL then
              break
          end if
      end for
      H^[j-1] ← HamiFun(T^[j-1], S^[j-1,k])                // coupling
      T^[j] ← ULMStep(T^[j-1], H^[j-1], Γ)
      if |T^[j] − T^[j-1]| < ε_UL then
          break
      end if
      S^[j,0] ← S^[j-1,k]
  end for
  return T^[j-1], S^[j-1,k], H^[j-1]
```

3.4. Convergence, stationarity and optimality

Let us discuss some properties of Algorithm 1. As discussed in Section 2, from a bilevel viewpoint, a solution to the LLP corresponds to a cost associated with a value of the scalar decision variable accounted for by the ULP. Hence, if both, the LLM and the ULM (globally) converge, possibly within prescribed tolerances, one could expect the overall algorithm to converge too. Some examples are in order. The Brent's method [14], the Newton's and the Halley's method [12, 13] could play the role of ULM, and usually these are guaranteed to converge under assumptions on the underlying function (h) and initial guess ($T^{[0]}$). For the LLM, the ASRE method globally converges under mild assumptions ($\mathbf{f}$ bounded and locally Lipschitz continuous) [9, Remark 4]; for the MASRE method, adopted in the following, global convergence is not guaranteed but expected [10]. Properties of the returned result may depend on the particular LLM and ULM. However, one cannot expect more than stationarity of the (reduced) cost with respect to the final time, because a necessary (and not sufficient) condition is exploited. Nevertheless, a sufficient condition can be *a posteriori* checked, see Section 3.1. This could guarantee local optimality. In general, in order to get a global optimum, both, the LLM and the ULM must guarantee global optimality (the vanishing condition cannot be exploited anymore, for it being only necessary in general).

Remark 3. *There might be a price for inexactness of the lower-level solution, possibly degrading not only the accuracy but the convergence properties too, because of the bilevel structure. On the other hand, however, one could opt for inexact lower-level solutions in order to speed up the execution, seeking a trade-off with the convergence behaviour and the desired accuracy, see Section 5.*

4. MASRE method

This section discusses how to exploit the MASRE, or *modified ASRE* method, proposed in [10], as a LLM and the coupling with the ULM. The ASRE method [9], or *approximating sequence of Riccati equation*, takes advantage of a state-dependent coefficient (SDC) factorization of the dynamics and introduces a sequence of time-varying linear-quadratic problems, which under mild assumptions globally converge to the original nonlinear problem [9]. Each of these subproblems is equivalent to a fixed finite time-horizon time-varying linear-quadratic regulator (TVLQR) and thus can be solved by classical tools, e.g. differential Riccati equation or state-transition matrix [9, 17, 10]. In fact, being the SDC factorization nonunique (for $n > 1$), different SDC matrices can be provided, but obtaining different results in terms of system trajectory and cost function; the MASRE aims at solving this issue while reducing the control effort and improving the convergence properties of the algorithm [10]. For example, given dynamics (1), a parametrized SDC factorization consists of two matrix-valued functions, $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, such that it holds $\boldsymbol{\alpha}(\mathbf{x},\mathbf{u},\mathbf{p})\mathbf{x} + \boldsymbol{\beta}(\mathbf{x},\mathbf{u},\mathbf{p})\mathbf{u} = \mathbf{f}(\mathbf{x},\mathbf{u})$ for any $\mathbf{x}$, $\mathbf{u}$ and $\mathbf{p}$. Then, in the k-th problem of the MASRE, the dynamics read $\dot{\mathbf{x}} = \boldsymbol{\alpha}^{[k-1]}\mathbf{x} + \boldsymbol{\beta}^{[k-1]}\mathbf{u}$, where $\boldsymbol{\alpha}^{[k-1]}$ and $\boldsymbol{\beta}^{[k-1]}$ are defined based on the solution of the previous problem and only depend on time. For a detailed discussion about this method, the proposed SDC selection strategy and the underlying assumptions on the original problem, see [9, 10]. The choice of this method stems from the fact that it allows to compute the derivatives of h_{LL} (8) with little overhead, and these can then be exploited by adopting a derivative-based ULM, e.g. the Halley's method [12, 13]. In the following, the TVLQR is briefly addressed, in order to prepare a common ground on which the coupling interface can easily be developed.

Time-Varying Linear-Quadratic Regulator. Let us consider a time interval $[0,T]$, with given final time $T > 0$, a linear time-varying state differential equation, coupled boundary conditions and a cost functional:

$$\dot{\mathbf{x}}(t) = \boldsymbol{A}(t)\mathbf{x}(t) + \boldsymbol{B}(t)\mathbf{u}(t) \quad \text{for a.e. } t \in [0,T] \tag{10}$$

$$\mathbf{c} = \boldsymbol{C}_0\mathbf{x}(0) + \boldsymbol{C}_T\mathbf{x}(T) \tag{11}$$

$$J(\mathbf{x},\mathbf{u}) := \int_0^T [\omega + \ell(\tau,\mathbf{x}(\tau),\mathbf{u}(\tau))]\,\mathrm{d}\tau \tag{12}$$

$$\ell(t,\mathbf{x},\mathbf{u}) := \frac{1}{2}\mathbf{x}^\top\boldsymbol{Q}(\tau)\mathbf{x} + \mathbf{u}^\top\boldsymbol{S}(\tau)\mathbf{x} + \frac{1}{2}\mathbf{u}^\top\boldsymbol{R}(\tau)\mathbf{u} \tag{13}$$

with given matrix-valued functions $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{Q}$, $\boldsymbol{S}$ and $\boldsymbol{R}$, matrices $\boldsymbol{C}_0, \boldsymbol{C}_T \in \mathbb{R}^{n_c \times n_x}$ and vector $\mathbf{c} \in \mathbb{R}^{n_c}$, obtained from a time-varying linear-quadratic approximation of the original problem [9, 10]. Then, it reads:

Problem 4 (TVLQR)**.** *Find a control $\mathbf{u} \in U(T)$ and a state $\mathbf{x} \in X(T)$ minimizing the cost functional J in (12) while satisfying the constraints (10)–(11).*

An approach based on the state transition matrix is adopted [17, 10, 8], exploiting the simple structure of the problem. Since the scheme has been already developed elsewhere, we just outline the derivation focusing on those features which are of interest for this work.

Recalling the Pontryagin's Minimum Principle and introducing multipliers $\boldsymbol{\lambda} : [0,T] \to \mathbb{R}^{n_x}$ and $\boldsymbol{\eta} \in \mathbb{R}^{n_c}$, one gets the following first-order necessary optimality conditions [1, 3], along with constraints (10)–(11):

$$\dot{\boldsymbol{\lambda}}(t) = -\boldsymbol{Q}(t)\mathbf{x}(t) - \boldsymbol{S}^\top(t)\mathbf{u}(t) - \boldsymbol{A}^\top(t)\boldsymbol{\lambda}(t) \quad \text{for a.e. } t \in [0,T] \tag{14a}$$

$$\mathbf{0} = \boldsymbol{S}(t)\mathbf{x}(t) + \boldsymbol{R}(t)\mathbf{u}(t) + \boldsymbol{B}^\top(t)\boldsymbol{\lambda}(t) \tag{14b}$$

$$-\boldsymbol{\lambda}(0) = \boldsymbol{C}_0^\top\boldsymbol{\eta} \tag{14c}$$

$$\boldsymbol{\lambda}(T) = \boldsymbol{C}_T^\top\boldsymbol{\eta} \tag{14d}$$

These conditions constitute a linear two-point boundary value problem (BVP). Collecting state and adjoint variables in $\mathbf{z} := (\mathbf{x},\boldsymbol{\lambda})$ and plugging (14b) into (10) and (14a), one obtains the Hamiltonian system $\dot{\mathbf{z}}(t) = \boldsymbol{M}(t)\mathbf{z}(t)$, where function $\boldsymbol{M} : [0,T] \to \mathbb{R}^{2n_x \times 2n_x}$ is built upon $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{Q}$, $\boldsymbol{S}$ and $\boldsymbol{R}$. Thanks to linearity and homogeneity, its solution reads $\mathbf{z}(t) = \boldsymbol{\varphi}(t)\mathbf{z}(0)$, where function $\boldsymbol{\varphi} : [0,T] \to \mathbb{R}^{2n_x \times 2n_x}$ denotes the state transition matrix [17], which, by definition, satisfies $\dot{\boldsymbol{\varphi}}(t) = \boldsymbol{M}(t)\boldsymbol{\varphi}(t)$ for all $t \in [0,T]$ and $\boldsymbol{\varphi}(0) = \boldsymbol{I}$, denoting $\boldsymbol{I}$ the identity matrix. Let us denote $\mathbf{z}_0 := \mathbf{z}(0)$ the initial condition. Then, rearranging boundary and transversality conditions, respectively (11) and (14d), one can obtain a linear system with $\mathbf{z}_0$ and $\boldsymbol{\eta}$ as unknowns. Collecting these variables in vector $\mathbf{s} := (\mathbf{z}_0, \boldsymbol{\eta})$ and highlighting the dependence on the final time T, the linear system reads

$$\boldsymbol{A}_s(T)\,\mathbf{s}(T) = \mathbf{b}_s\ , \tag{15}$$

where function $\boldsymbol{A}_s : \Gamma \to \mathbb{R}^{n_s \times n_s}$, $n_s = 2n_x + n_c$, depends on $\boldsymbol{\varphi}$, $\boldsymbol{C}_0$ and $\boldsymbol{C}_T$, while vector $\mathbf{b}_s \in \mathbb{R}^{n_s}$ on $\mathbf{c}$ only. Given final time T, the (unique) solution $\mathbf{s}^\star$ to linear system (15) contains the (unique) initial conditions $\mathbf{z}_0^\star$ for the Hamilton system from which the (unique globally optimal) state $\mathbf{x}^\star$ can be reconstructed, along with the corresponding optimal control $\mathbf{u}^\star$ and the adjoint $\boldsymbol{\lambda}^\star$, for the given final time T.

Root-finding procedure. The Hamilton function associated with the TVLQR is defined as

$$\mathcal{H}_{\mathrm{LL}}(t, \mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}) := \omega + \ell(t, \mathbf{x}, \mathbf{u}) + \boldsymbol{\lambda}^\top [\boldsymbol{A}(t)\mathbf{x} + \boldsymbol{B}(t)\mathbf{u}] \quad (16)$$

Along a solution, the optimal control satisfies (14b), hence, under suitable assumptions, it is possible to express $\mathbf{u}_T$ in feedback form, namely

$$\mathbf{u}_{\mathrm{fb}}(t, \mathbf{x}, \boldsymbol{\lambda}) := -\boldsymbol{R}^{-1}(t) \left[\boldsymbol{S}(t)\mathbf{x} + \boldsymbol{B}^\top(t)\boldsymbol{\lambda}\right] \quad (17)$$

Then, for any given final time T, from (5), (7) and (8), choosing $\tau = 0$, we get

$$h(T) = \mathcal{H}_{\mathrm{LL}}\left(0, \mathbf{x}_T(0), \mathbf{u}_{\mathrm{fb}}(0, \mathbf{x}_T(0), \boldsymbol{\lambda}_T(0)), \boldsymbol{\lambda}_T(0)\right) \quad (18)$$

and one can notice that $\mathbf{x}_T(0)$ and $\boldsymbol{\lambda}_T(0)$ are the initial conditions for the Hamilton system obtained by solving the linear system (15) associated with the TVLQR of the corresponding MASRE iteration, for the given final time T. Let us denote $\mathbf{z}_0(T)$ the vector of initial conditions obtained for final time T. From (13) and (16)–(18), with some rearrangements, one can write

$$h(T) = \omega + \frac{1}{2}\mathbf{z}_0^\top(T)\, \boldsymbol{W}_0\, \mathbf{z}_0(T) \quad (19)$$

where $\boldsymbol{W}_0$ is a symmetric matrix obtained by rearranging the four $n_x \times n_x$ blocks of the Hamilton matrix $\boldsymbol{M}(0)$, defined above. Notice that $\boldsymbol{W}_0$ does not depend explicitly on the final time T but implicitly through the underlying linear time-varying approximation. Based on (19), one can obtain the derivatives of h and exploit them within the ULM.

$$h'(T) = \mathbf{z}_0^\top(T)\boldsymbol{W}_0\mathbf{z}_0'(T) \quad (20a)$$

$$h''(T) = \mathbf{z}_0^\top(T)\boldsymbol{W}_0\mathbf{z}_0''(T) + (\mathbf{z}_0')^\top(T)\boldsymbol{W}_0\mathbf{z}_0'(T) \quad (20b)$$

Herein, vectors $\mathbf{z}_0'(T)$ and $\mathbf{z}_0''(T)$ are the derivatives of $\mathbf{z}_0$ at T and can be evaluated starting from linear system (15), by solving two additional linear systems in sequence; namely

$$\boldsymbol{A}_s(T)\mathbf{s}'(T) = -\boldsymbol{A}_s'(T)\mathbf{s}(T)$$
$$\boldsymbol{A}_s(T)\mathbf{s}''(T) = -\boldsymbol{A}_s''(T)\mathbf{s}(T) - 2\boldsymbol{A}_s'(T)\mathbf{s}'(T)$$

with $\boldsymbol{A}_s'(T)$ and $\boldsymbol{A}_s''(T)$ derivatives of $\boldsymbol{A}_s$ at T. The time derivative of the state transition matrix $\boldsymbol{\varphi}$ comes from its definition, and the sensitivities of $\boldsymbol{\varphi}$ to T at time T are $\boldsymbol{\varphi}'(T) = \boldsymbol{M}(T)\boldsymbol{\varphi}(T)$ and $\boldsymbol{\varphi}''(T) = [\dot{\boldsymbol{M}}(T) + \boldsymbol{M}^2(T)]\boldsymbol{\varphi}(T)$, being $\boldsymbol{M}$ the Hamilton matrix, whose time derivative depends on that of $\boldsymbol{A}$, $\mathbf{B}$, $\mathbf{Q}$, $\mathbf{S}$ and $\mathbf{R}$. These, in turn, are related to the underlying trajectory and factorization.

Remark 4. *We stress that (implicit) sensitivity of the Hamiltonian with respect to the final time might differ, between TVLQR and original nonlinear problem, since the LTV approximation likely neglects some dependencies. In fact, matrix $\boldsymbol{W}_0$ implicitly depends on the final time, but not explicitly in* (19)–(20).

5. Numerical results

5.1. Problem

We consider the trajectory planning of a point-to-point manoeuvre for a vehicle, with a simplified single-track model consisting of control variables $\mathbf{u} = (a, \omega)$ (acceleration and steering velocity), state variables $\mathbf{x} = (x, y, \theta, v, \delta)$ (coordinates of the centre of gravity in a fixed global reference frame, yaw angle, steering angle and speed) and parameters $\boldsymbol{\pi} = (L, c_1, c_2)$ (wheelbase, linear and quadratic air drag coefficients). The equations of motion are given in the form of (1), with $\mathbf{f}$ defined by

$$\mathbf{f}(\mathbf{x}, \mathbf{u}) := \begin{pmatrix} v\cos\delta\cos\theta \\ v\cos\delta\sin\theta \\ v\sin\delta/L \\ a - c_1 v - c_2 v^2 \\ \omega \end{pmatrix}. \quad (21)$$

One can verify that $\mathbf{f}$ satisfies the assumptions for the MASRE to work [9]. Appearing the control linearly in (21), a suitable constant matrix $\boldsymbol{\beta}$ is defined and a parametric decomposition $\boldsymbol{\alpha}(\mathbf{x}, \mathbf{p})$ of $\mathbf{f}(\mathbf{x}, \mathbf{u})$ is needed. In order to keep low the dimension of the SDC parametrization, a factorization with only 2 parameters is chosen; it reads:

$$\boldsymbol{\alpha}(\mathbf{x}, \mathbf{p}) := \begin{bmatrix} 0 & 0 & 0 & \cos\delta\cos\theta & 0 \\ 0 & 0 & (1-p_1)v\frac{\sin\theta}{\theta}\cos\delta & p_1\cos\delta\sin\theta & 0 \\ 0 & 0 & 0 & (1-p_2)\frac{\sin\delta}{L} & p_2\frac{v}{L}\frac{\sin\delta}{\delta} \\ 0 & 0 & 0 & -c_1 - c_2 v & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

where $\mathbf{p} = (p_1, p_2) \in [0, 1]^2$ and it satisfies the property reported in Section 4. Furthermore, the cost functional mirrors that defined in (12)–(13), with time-invariant weighting matrices; also boundary conditions are given in the form of (11). For the numerical computations we used the following parameters: $\boldsymbol{\pi} = (4, 10^{-3}, 10^{-4})$, $\boldsymbol{R} = \mathrm{diag}(1, 10)$, $\boldsymbol{Q} = \mathrm{diag}(0, 0, 0, 0, 10)$, $\boldsymbol{S} = \mathbf{0}$, $\omega = 10$. Boundary conditions are embedded in vectors $\boldsymbol{\xi}_0 = (0, 0, 0, 10, 0)$ and $\boldsymbol{\xi}_T = (100, 20, 0, 10, 0)$. Given a final time guess $T_{[0]} > 0$, the state guess is defined as $\mathbf{x}_{[0]}(t) := \boldsymbol{\xi}_0 + (\boldsymbol{\xi}_T - \boldsymbol{\xi}_0)t/T_{[0]}$ for any $t \in [0, T_{[0]}]$. The final time is in the box $\Gamma = [5, 20]$. The initial SDC parametrization is $\mathbf{p}_{[0]} = (0.5, 0.5)$ and several inital guesses for the final time have been tested, with $T_{[0]} \in [5, 15]$. In the following, two examples are considered, both with fixed initial state, namely $\mathbf{x}(0) = \boldsymbol{\xi}_0$. The first example (E$\theta$) imposes a given final position and yaw angle; similarly, the second (Ev) prescribes a given final position and velocity. Matrices $\boldsymbol{C}_0$, $\boldsymbol{C}_T$ and vector $\mathbf{c}$ in (11) are built accordingly.

The parameters controlling the execution of Algorithm 1 are as follows. MASRE: $\epsilon_{\text{LL}}^{\text{abs}} = 10^{-8}$, $\epsilon_{\text{LL}}^{\text{rel}} = 10^{-6}$, $N_{\text{LL}} = 1$, $N_{\text{SDC}} = 10$ (for the SDC selection). Halley's method: $\epsilon_{\text{UL}}^{\text{abs}} = 10^{-6}$, $\epsilon_{\text{UL}}^{\text{rel}} = 10^{-5}$, $N_{\text{UL}} = 20$.

5.2. Results

Figures 1–2 report the evolution of the most interesting quantities during execution of Algorithm 1, for different initial guesses, while Figures 4–3 display state and control generated at different stages of the solution process, for a given initial guess. After 6-8 iterations, final time and cost function stabilize already, Fig. 1. For the two examples, each of these quantities converges to a single value, independent on the initial guess. The value of h approaches zero relatively quickly, while h' attains positive values, Fig. 1. The SDC parametrization moves away from the initial guess but does not always settle to a steady optimal value, Fig. 2. The planned path is very close to the final solution after just four iterations, even though the terminal conditions (with the nonlinear dynamics) are not satisfied yet, Fig. 4. A similar observation holds for the other states and controls, Fig. 3.

For each example, convergence toward a single value is achieved for the optimal final time, regardless the initial guess, Fig. 1. Reached final times correspond to a minimum; notice $h \approx 0$ (necessary) and $h' > 0$ (sufficient) in Fig. 1. A number of local minima (and maxima) might be expected, but this is likely not the case for the problem faced here. Notice also that SDC parametrization continues to change, Fig. 2, even if state and control trajectories stabilize, Fig. 4–3. The number of lower-level iterations per each upper-level iteration has been limited; in fact, N_{LL} and N_{SDC} have purposely small values. This means that, especially during the first iterations, the LLM returns inexact, or at least inaccurate, solutions to the LLP. Consequently, the reduced Hamiltonian and its sensitivity are inaccurately computed. However, in spite of this inexactness and Remark 3, the algorithm generates a sequence of final times that converges in relatively few iterations.

The developed method and the testing environment for these numerical tests have been programmed in plain MATLAB language. With this implementation, the computational effort in terms of elapsed time for a MASRE iteration is in the order of a second, while the root-finding procedure, consisting of Hamilton function evaluation and Halley's step, takes less than a millisecond. Thus, the overhead required for the final time optimization is negligible compared to the solution for the fixed-time problem.

For the sake of comparison with a standard approach, the OCP was also solved, through time transformation, using the software package OCPID-DAE1 [15] (direct single shooting, 101 equidistant grid points, continuous piecewise linear control parametrization), starting from a relatively good initial guess. The solution approximately matched

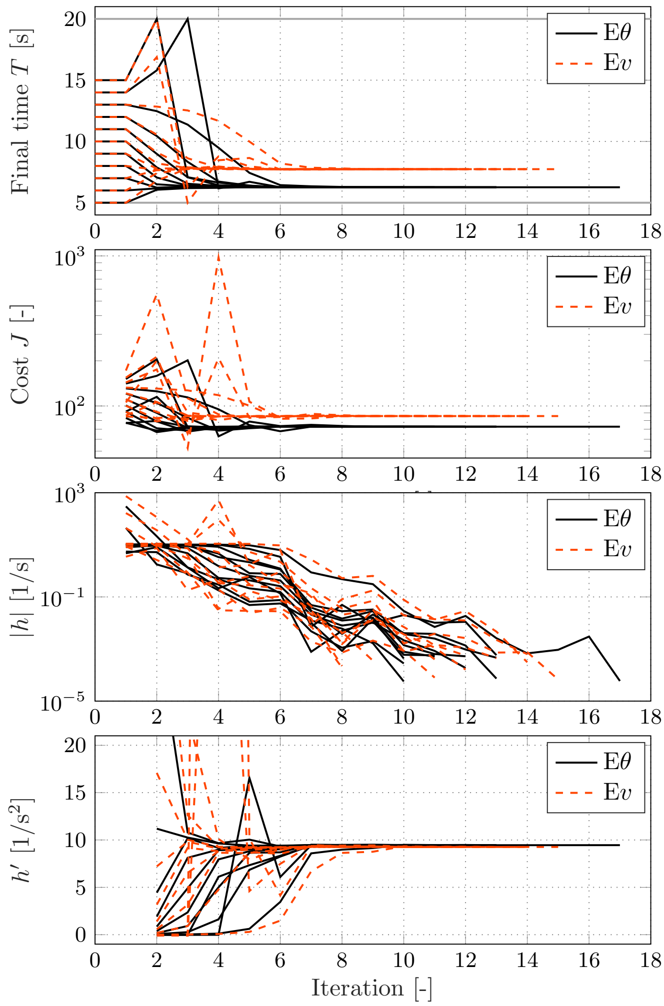

Figure 1: Final time T, cost J, Hamiltonian h and its derivative h' during algorithm iterations, for different initial guesses $T_{[0]}$.

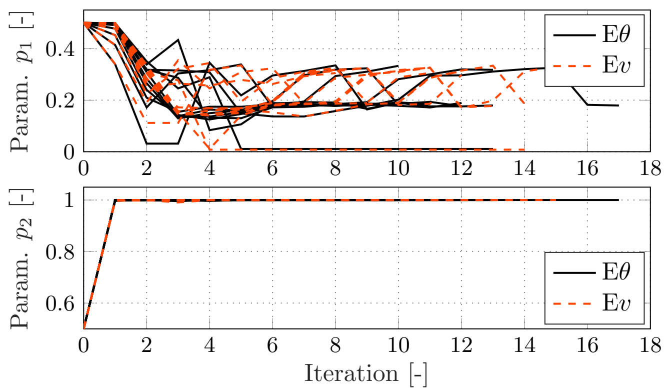

Figure 2: SDC parameters $\mathbf{p}$ during algorithm iterations, for different initial guesses $T_{[0]}$.

the one obtained previously, but with far more iterations and difficulties in convergence.

6. Conclusion

This paper introduced a bilevel approach to free final time nonlinear optimal control problems and a numerical method to approximate their solution. The optimization of final time and control input are almost decoupled, requiring the two levels to exchange only information about the Hamilton function. It remains to be seen how Algorithm 1 performs compared to direct methods with a time transformation. Moreover, it would be interesting to adopt other direct methods to solve the fixed-time problem and to investigate derivative-free algorithms to tackle the upper level.

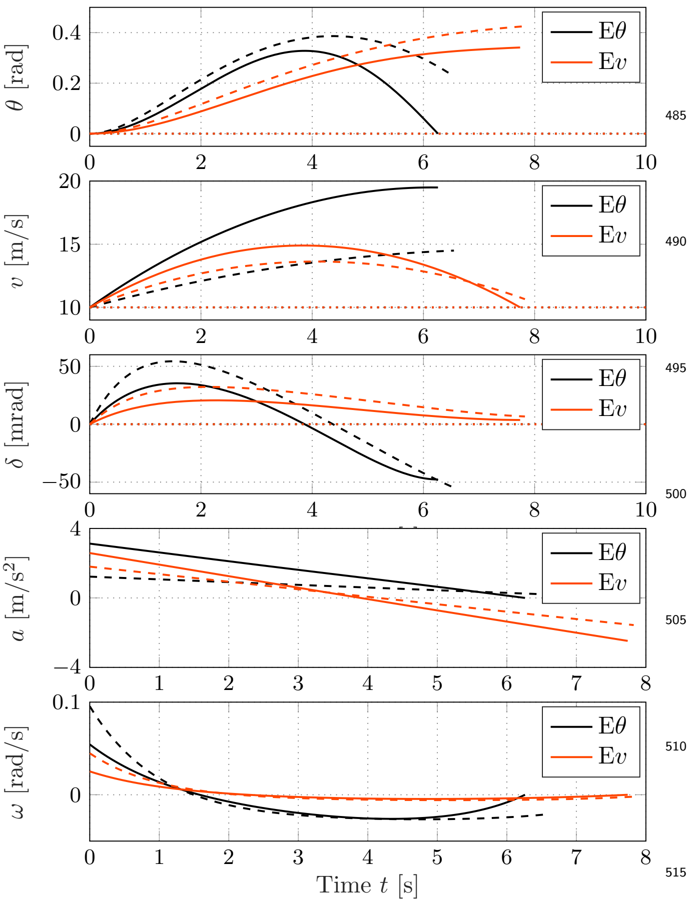

Figure 3: State (attitude, speed and steering angle) and control (acceleration and steering velocity), with initial guess $T_{[0]} = 10$: first (dotted), fourth (dashed) and last (solid) iteration.

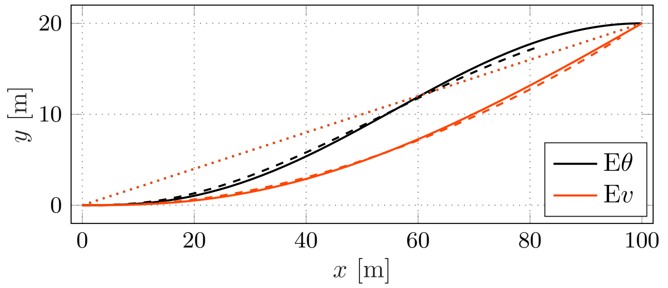

Figure 4: Vehicle path, with initial guess $T_{[0]} = 10$: first (dotted), fourth (dashed) and last (solid) iteration.

References

[1] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, E. Mishchenko, The mathematical theory of optimal processes (International series of monographs in pure and applied mathematics, Interscience Publishers, 1962.

[2] A. Locatelli, Optimal control: An introduction., Basel: Birkhäuser, 2001.

[3] M. Gerdts, Optimal Control of ODEs and DAEs, De Gruyter, 2011. doi:10.1515/9783110249996.

[4] P. Bosetti, F. Biral, Application of optimal control theory to milling process, in: IECON 2014 - 40th Annual Conference of the IEEE Industrial Electronics Society, 2014, pp. 4896–4901. doi:10.1109/IECON.2014.7049243.

[5] R. Lot, F. Biral, A curvilinear abscissa approach for the lap time optimization of racing vehicles, IFAC Proceedings Volumes 47 (3) (2014) 7559–7565, 19th IFAC World Congress. doi:10.3182/20140824-6-ZA-1003.00868.

[6] L. Van den Broeck, M. Diehl, J. Swevers, Time optimal MPC for mechatronic applications, in: Proceedings of the IEEE Conference on Decision and Control (CDC), Shanghai, China, 2009, pp. 8040–8045. doi:10.1109/CDC.2009.5400667.

[7] K. D. Palagachev, M. Gerdts, Exploitation of the value function in a bilevel optimal control problem, in: L. Bociu, J.-A. Désidéri, A. Habbal (Eds.), System Modeling and Optimization, Vol. 494, Springer, Cham, 2016, pp. 410–419, CSMO 2015, IFIP. doi:10.1007/978-3-319-55795-3_39.

[8] A. De Marchi, M. Gerdts, Free finite horizon LQR: a bilevel perspective and its application to model predictive control, Automatica (2018 in press) 1–16.

[9] T. Çimen, S. P. Banks, Global optimal feedback control for general non-linear system with non-quadratic performance criteria, System & Control Letters 53 (5) (2004) 327–346. doi:10.1016/j.sysconle.2004.05.008.

[10] F. Topputo, M. Miani, F. Bernelli-Zazzera, Optimal selection of the coefficient matrix in state-dependent control methods, Journal of Guidance Control and Dynamics 38 (2015) 861–873. doi:10.2514/1.G000136.

[11] L. Grüne, J. Pannek, Nonlinear model predictive control. Theory and algorithms., London: Springer, 2011. doi:10.1007/978-3-319-46024-6.

[12] A. S. Householder, The Numerical Treatment of a Single Nonlinear Equation, International series in pure and applied mathematics, New York: McGraw-Hill, 1970.

[13] G. Alefeld, On the convergence of Halley's method, The American Mathematical Monthly 88 (7) (1981) 530–536.

[14] R. P. Brent, An algorithm with guaranteed convergence for finding a zero of a function, Computer Journal 14 (4) (1971) 422–425.

[15] M. Gerdts, OCPID-DAE1 — optimal control and parameter identification with differential-algebraic equations of index 1, Tech. rep., University of the Federal Armed Forces at Munich, User's Guide, http://www.optimal-control.de (2013).

[16] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, M. Diehl, CasADi: a software framework for nonlinear optimization and optimal control, Mathematical Programming Computationdoi:10.1007/s12532-018-0139-4.

[17] J. P. Hespanha, Linear Systems Theory, Princeton Press, Princeton, New Jersey, 2009, iSBN13: 978-0-691-14021-6.