

acados – a modular open-source framework for fast embedded optimal control

Robin Verschueren · **Gianluca Frison** ·
Dimitris Kouzoupis · **Jonathan Frey** ·
Niels van Duijkeren · **Andrea Zanelli** ·
Branimir Novoselnik · **Thivaharan**
Albin · **Rien Quirynen** · **Moritz Diehl**

Manuscript as accepted by Mathematical Programming Computation

Abstract This paper presents the `acados` software package, a collection of solvers for fast embedded optimization intended for fast embedded applications. Its interfaces to higher-level languages make it useful for quickly designing an optimization-based control algorithm by putting together different algorithmic components that can be readily connected and interchanged. Since the core of `acados` is written on top of a high-performance linear algebra li-

This research was supported by the EU via ERC-HIGHWIND (259 166), FP7-ITN-TEMPO (607 957), H2020-ITN-AWESCO (642 682), by the DFG via the Research Unit FOR 2401 and project 424107692, and by the Federal Ministry for Economic Affairs and Energy (BMW) via `eco4wind` (0324125B) and `DyConPV` (0324166B).

Robin Verschueren
ABB Corporate Research, Baden-Dättwil, Switzerland

Gianluca Frison
Dimitris Kouzoupis
Andrea Zanelli
IMTEK, University of Freiburg, Germany

Niels van Duijkeren
Department of Mechanical Engineering, KU Leuven, Belgium

Branimir Novoselnik
Faculty of Electrical Engineering and Computing, University of Zagreb, Croatia

Thivaharan Albin
Institute for Dynamic Systems and Control, ETH Zürich, Switzerland

Rien Quirynen
Mitsubishi Electric Research Labs, Cambridge, MA, USA

Jonathan Frey
Moritz Diehl
IMTEK and Faculty of Mathematics, University of Freiburg, Germany

Note: the first four authors contributed an equal amount to the work presented in this article.

brary, we do not sacrifice computational performance. Thus, we aim to provide both flexibility and performance through modularity, without the need to rely on automatic code generation, which facilitates maintainability and extensibility. The main features of `acados` are: efficient optimal control algorithms targeting embedded devices implemented in C, linear algebra based on the high-performance `BLASFEO` [35] library, user-friendly interfaces to `MATLAB` and `Python`, and compatibility with the modeling language of `CasADi` [4]. `acados` is free and open-source software released under the permissive BSD 2-Clause license.

Keywords Direct optimal control · Optimization algorithms

Mathematics Subject Classification (2000) 49-04

1 Introduction

Embedded optimization, according to the definition in [27], is solving optimization problems *autonomously* and *with limited resources*. The topic of this article is embedded optimal control, an important class of methods within embedded optimization. It focuses on calculating optimal decisions in order to control a dynamic system as the state changes. There exist many algorithms for embedded optimal control and quite a few were successfully applied to real-time and embedded applications such as robotic trajectory optimization [83], autonomous driving [62] and drones [46]. For an overview on embedded optimization, we refer the reader to [27]. One of the most popular approaches in embedded optimal control nowadays is model predictive control (MPC) [65, 81, 43]. It is based on predicting the future behavior of a system and using this information to optimize for the action at the current time step. In linear MPC (also called linear-quadratic MPC (LQMPC)), the constraints, including the dynamic model, are affine and the objective is quadratic. With nonlinear MPC (NMPC), some or all of the constraints and objective are nonlinear functions¹. A related problem is that of moving horizon estimation (MHE) for estimating states and parameters online. We will present a unified notation of NMPC and MHE problems in Section 2.

Historically, MPC was primarily applied to systems with long timescales, most notably in chemical processing [74], due to the fact that during each time step, a computationally costly optimization problem has to be solved. More recent algorithmic developments [21, 61, 71, 9] and increasingly powerful embedded hardware render MPC real-time feasible for applications with shorter timescales such as autonomous driving, robotics, and avionics, as in some of the references cited above. A big role in bringing MPC to real-time applications is played by the implementation of efficient embedded optimal control methods, giving rise to software packages such as `MPT` [47] for explicit MPC, `qpOASES` [28], an active-set solver for quadratic programming

¹ In this paper, we focus on finite-dimensional problems, i.e., the objective is not a functional.

(QP), FORCES [24,96], an interior-point solver for quadratically constrained quadratic programming (QCQP) and nonlinear programs with optimal control structure, and the ACADO Code Generation tool [49] for tailored sequential quadratic programming (SQP) based NMPC solvers. Other examples of nonlinear embedded optimization packages are VIATOC [50], GRAMPC [26] and FalcoOpt [88], to which we compare `acados` in Section 5. A non-exhaustive list of embedded optimization software packages can be found in Table 1.

Ultimately, embedded software should run on a dedicated hardware platform. In this paper, we focus solely on central processing units (CPU) as the algorithms presented typically don't profit as much from the parallelization capabilities of massively parallel hardware platforms such as graphical processing units (GPU), tensor processing units (TPU) or field programmable gate arrays (FPGA). Please note that we don't pin ourselves down to a specific CPU architecture: `acados` has been shown to work with x86, x86_64, ARMv7A, ARMv8A and PowerPC architectures.

One challenge in developing software for embedded optimal control lies in the trade-off between flexibility, memory usage and speed. Many of the software packages mentioned in Table 1 are based on *automatic code generation*. One reason for that is to have self-contained efficient linear algebra routines. Often however, the size of the problem and the choice of algorithms are then fixed for one specific optimal control instance, inducing a loss of flexibility. In some cases, a compiler or a human could generate faster or more memory-efficient code. For example, for linear algebra operations, the recently developed high-performance linear algebra package BLASFEO [35] outperforms code-generated triple-loop linear algebra routines and state of the art BLAS and LAPACK implementations for the moderate matrix size typical for embedded optimization applications. Since the linear algebra operations typically comprise most of the computational complexity of the algorithms considered, we can use standard compiler optimizations for the code surrounding the linear algebra routines without noticeably sacrificing performance (see Section 5). By basing `acados` on BLASFEO, we exchange self-containedness for performance and flexibility. We believe this is a better trade-off.

Another challenge for embedded optimal control software is related to the process of software development. Often, to not sacrifice speed of execution and/or memory footprint, embedded optimal control software uses global data and suffers from tight coupling between algorithmic components. This might lead to a codebase that is difficult to understand, maintain and extend. We choose, as opposed to some other embedded optimal control software packages, to avoid these pitfalls by not unnecessarily sacrificing maintainability and readability of the codebase for a small gain in performance and/or a reduction of memory footprint. We try to achieve this goal by organizing the code in a modular fashion, with formal interfaces between the different algorithmic components, as described in Section 4. This allows for a straightforward way of interchanging solvers, routines, and libraries needed for the embedded control algorithm.

Table 1 SOFTWARE PACKAGES FOR EMBEDDED OPTIMIZATION.

Software	Year	Latest Reference	Targets	License
MUSCOD-II	1997	[60]	NMPC	proprietary
HQP	1997	[29]	NLP	LGPL v2
MPC Toolbox	1998	[66]	LQMPC	proprietary
AutoGenU	2000	[71]	NMPC	proprietary
OOQP	2001	[38]	QP	proprietary
MPT3	2003	[47]	expl. MPC	GPL
Hybrid toolbox	2003	[8]	expl. MPC	proprietary
qpOASES	2006	[28]	QP	LGPL v2.1
PQP	2008	[20]	QP	proprietary
CVXGEN	2009	[67]	LP, QP	proprietary
ACADO Codegen	2009	[49]	NMPC	LGPL v3
FiOrdOs	2011	[89]	QP	GPL v3
muAO-MPC	2013	[99]	LQMPC	3-clause BSD
FORCES	2011	[24]	QP, QCQP	proprietary
ECOS	2013	[23]	SOCP	GPL v3
GRAMPC 1.0	2014	[51]	NMPC	LGPL v3
qpDUNES	2014	[30]	LQMPC	LGPL v3
DuQuad	2014	[58]	QP	unknown
HPMPC	2014	[37]	LQMPC	LGPL v2.1
VIATOC	2015	[50]	NMPC	GPL v3 and LGPL v3
PIPS-NLP	2016	[17]	NLP	3-clause BSD
GPAD [†]	2016	[73]	QP	GPL v3
Forces NLP	2017	[96]	NMPC	proprietary
OSQP	2017	[87]	QP	Apache v2.0
FalcOpt	2017	[88]	NMPC	MIT
HPIPM	2017	[34]	LQMPC, QP	2-clause BSD
ODYS	2017	[18]	QP	proprietary
protoip	2017	[53]	NMPC	unknown
SPLIT toolbox	2017	[84]	LQMPC	unknown
Control Toolbox	2018	[39]	NMPC	2-clause BSD
GRAMPC 2.0	2018	[26]	NMPC	LGPL v3
PANOC.jl	2018	[82]	NMPC	unknown
nmpc-codegen	2018	[68]	NMPC	LGPL v3
PRESAS	2018	[79]	QP	proprietary
ParNMPC	2018	[19]	NMPC	unknown
qrqp	2018	[5]	QP	LGPL v3
acados	2018	[91]	NMPC	2-clause BSD
TMPC	2018	[52]	NMPC	unknown
qpSWIFT	2019	[72]	QP	unknown
MATMPC	2019	[15]	NMPC	GPL v3
OpEn	2019	[85]	NMPC	Apache v2.0
PolyMPC	2019	[63]	NMPC	MIT
QPALM	2020	[48]	QP	GPL v3
NASOQ	2020	[16]	QP	MIT
ASIPM	2020	[31]	QP	proprietary

[†] As implemented by Pantelis Sopasakis.

A final aspect of embedded optimal control software that affects flexibility, memory and runtime is the choice of modeling language and corresponding automatic differentiation tool. Several modeling languages exist, e.g., `Mathematica`, `sympy` or the `MATLAB Symbolic Toolbox`. Many of these languages make use of expression trees to represent mathematical functions, which potentially leads to a large code size, high memory usage and slow evaluation of higher-order derivatives for non-trivial models. On the contrary, the `CasADi` [4] modeling language is based on expression *graphs*. This often leads to shorter instruction sequences and to smaller, typically faster code, which makes it more suitable for embedded applications. Also, it is free and open-source software. For these reasons, we favor `CasADi` for modeling nonlinear functions and differential-algebraic equations. Additionally, `acados` supports the use of hand-written or code-generated dynamic models as C source files.

In summary, the contribution of this paper is a new software package for embedded optimal control, called `acados`. It offers the following main features:

- efficient optimal control algorithms implemented in C,
- modular architecture enabling rapid prototyping of solution algorithms,
- interfaces to `Python` and `MATLAB`,
- high-performance linear algebra based on `BLASFEO` [35],
- compatible with `CasADi` expressions [4],
- deployable on a variety of embedded devices,
- publicly available as permissively licensed free and open-source software.

The remainder of the paper is organized as follows: in Section 2, we review some important elements of nonlinear embedded optimization algorithms relevant to `acados`. We discuss recent advances in embedded optimization algorithms that motivate the development of `acados` in Section 3. The software package `acados` itself is introduced in Section 4. Various numerical experiments including hardware-in-the-loop simulations as well as comparisons to other embedded optimization packages are presented in Section 5, and the paper is concluded in Section 6.

2 Algorithmic ingredients for embedded nonlinear optimal control

In this section, we first introduce the problems that typically have to be solved for NMPC and MHE. We introduce a general formulation, that can facilitate both multiple-shooting discretized NMPC and MHE problems. Subsequently, we describe the algorithmic ingredients of SQP-type methods as implemented in `acados` for the general problem.

2.1 Nonlinear Optimal Control

Let us first regard the continuous-time nonlinear optimal control problem (OCP) of the form

$$\underset{x(\cdot), z(\cdot), u(\cdot)}{\text{minimize}} \quad \int_0^T \ell(x(t), z(t), u(t)) dt + M(x(T)) \quad (1a)$$

$$\text{subject to} \quad x(0) = \bar{x}_0, \quad (1b)$$

$$0 = f(\dot{x}(t), x(t), z(t), u(t)), \quad t \in [0, T], \quad (1c)$$

$$0 \geq g(x(t), z(t), u(t)), \quad t \in [0, T]. \quad (1d)$$

In this notation, $x : \mathbb{R} \rightarrow \mathbb{R}^{n_x}$ denotes the *differential states*, $z : \mathbb{R} \rightarrow \mathbb{R}^{n_z}$ are the *algebraic variables* and $u : \mathbb{R} \rightarrow \mathbb{R}^{n_u}$ denotes the *control inputs*. Furthermore, we use $\ell : \mathbb{R}^{n_x} \times \mathbb{R}^{n_z} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}$ for the Lagrange term or running cost, $M : \mathbb{R}^{n_x} \rightarrow \mathbb{R}$ for the Mayer term or terminal cost. The dynamics are modeled with a set of implicit differential-algebraic equations (DAE) with right-hand-side $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_z} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x} \times \mathbb{R}^{n_z}$. In the remainder of the paper, we assume the implicit DAE to be of index 1, i.e. $\partial f / (\partial \dot{x}, \partial z)$ is invertible. The nonlinear path constraints are given by $g : \mathbb{R}^{n_x} \times \mathbb{R}^{n_z} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_g}$, and the initial value of the states is $\bar{x}_0 \in \mathbb{R}^{n_x}$. We consider the horizon length T to be fixed.

2.2 Multiple Shooting Discretization

In *acados*, we discretize nonlinear OCPs with a *multiple shooting* approach [13]. We introduce a time grid $[t_0, t_1, \dots, t_N]$ with $t_k < t_{k+1}, k = 0, \dots, N-1$, discrete state variables x_0, \dots, x_N , algebraic variables z_0, \dots, z_{N-1} and controls u_0, \dots, u_{N-1} . For the control trajectory, we choose a piecewise constant control parametrization. On each time interval $[t_k, t_{k+1})$, we can then write the result of the numerical simulation routine as

$$\begin{bmatrix} x_{k+1} \\ z_k \end{bmatrix} = \phi_k(x_k, u_k), \quad k = 0, \dots, N-1.$$

where the separate components will be denoted by $\phi_k^x(x_k, u_k)$ and $\phi_k^z(x_k, u_k)$.

The multiple shooting approach can often lead to better convergence behavior compared to single shooting, where the simulation and optimization is performed sequentially, as shown in [10]. The resulting nonlinear programming

(NLP) formulation looks as follows:

$$\underset{\substack{x_0, \dots, x_N, \\ z_0, \dots, z_{N-1}, \\ u_0, \dots, u_{N-1}}}{\text{minimize}} \quad \sum_{k=0}^{N-1} (t_{k+1} - t_k) \cdot \ell(x_k, z_k, u_k) + M(x_N) \quad (2a)$$

$$\text{subject to} \quad x_0 = \bar{x}_0, \quad (2b)$$

$$\begin{bmatrix} x_{k+1} \\ z_k \end{bmatrix} = \phi_k(x_k, u_k), \quad k = 0, \dots, N-1, \quad (2c)$$

$$0 \geq g_k(x_k, z_k, u_k) \quad k = 0, \dots, N-1. \quad (2d)$$

The optimal control formulation above is not the most general one that `acados` can handle: among others moving horizon estimation (MHE) problems, constraint relaxation via slack variables and equality constraints are supported as well. The next section presents a more general class of optimization problems which is handled by `acados`.

2.3 General Nonlinear Optimal Control structured Optimization Problem

In order to facilitate many different OCP formulations occurring in practice, `acados` uses the following general formulation of nonlinear optimal control structured optimization problems:

$$\underset{\substack{x_0, \dots, x_N, \\ u_0, \dots, u_{N-1}, \\ z_0, \dots, z_{N-1}, \\ s_0, \dots, s_N}}{\text{minimize}} \quad \sum_{k=0}^{N-1} l_k(x_k, u_k, z_k) + M(x_N) + \sum_{k=0}^N \rho_k(s_k) \quad (3a)$$

$$\text{subject to} \quad \begin{bmatrix} x_{k+1} \\ z_k \end{bmatrix} = \phi_k(x_k, u_k), \quad k = 0, \dots, N-1, \quad (3b)$$

$$0 \geq g_k(x_k, z_k, u_k) - J_{s,k} s_k \quad k = 0, \dots, N-1, \quad (3c)$$

$$0 \geq g_N(x_N) - J_{s,N} s_N, \quad (3d)$$

$$0 \leq s_k \quad k = 0, \dots, N. \quad (3e)$$

Here, the initial state constraint from (1b) is contained in the inequality constraint (3c) for $k = 0$.

Note that this general formulation also includes slack variables s_k which could alternatively be formulated as control inputs u_k . However, the slack variables s_k do not enter the system dynamics (3b) and can only enter the cost linearly and quadratically, i.e. via the functions $\rho_k(\cdot)$ of the form

$$\rho_k(s_k) = \sum_{i=1}^{n_{s_k}} \alpha_k^i s_k^i + \beta_k^i s_k^{i^2},$$

with $\alpha_k^i \in \mathbb{R}$, $\beta_k^i > 0$. These properties motivate the separation between slack and control variables and allow for an efficient treatment in solution methods as implemented in `acados`. The slack variables can be used to formulate soft

constraints or model piecewise quadratic possibly asymmetric cost functions among others. Soft constraints are often useful in practice, for example to deal with constraint violations due to model-plant mismatch or disturbances that would yield infeasibility of (2).

Moreover, `acados` is able to handle the most common types of cost and constraint functions in a tailored fashion. Regarding the cost functions $l_k(\cdot)$, $M(\cdot)$, `acados` is capable of exploiting the structure of the widely used linear and nonlinear least-squares functions, while also being able to handle general nonlinear functions.

Within the constraint functions $g_k(\cdot)$, `acados` is able to exploit the most common constraint types, which are simply bounds on x_k and u_k , linear constraints in x_k and u_k , such that only truly nonlinear constraints have to be formulated and treated as such, and one could write $g_k(\cdot)$ as

$$g_k(x_k, z_k, u_k) = \begin{bmatrix} J_{\text{bx},k}x_k \\ J_{\text{bu},k}u_k \\ C_{\text{x},k}x_k + C_{\text{u},k}u_k \\ g_{\text{nonl}}(x_k, z_k, u_k) \end{bmatrix}, \quad (4)$$

where $J_{\text{bx},k}$, $J_{\text{bu},k}$ are selection matrices that determine which components of x_k , u_k have simple bounds. Additionally, it is common to have upper and lower bounds on a constraint component of g_k , which allows for a more efficient treatment of these constraints within `acados`.

The above NLP (3) could be solved by any general-purpose NLP solver, like IPOPT [94]. The current scope of `acados`, however, encompasses efficient *embedded* optimal control methods for solving such structured NLPs, since these are better suited in a real-time and/or embedded setting [22]. Sequential Quadratic Programming (SQP) is an example of an embeddable optimization method and will be discussed in the following section.

2.4 Sequential Quadratic Programming and Real-Time Iterations

We briefly present the structure of an SQP algorithm as implemented in `acados`. At the very least, an embedded SQP algorithm should feature:

- numerical integration of the continuous-time dynamics,
- generation of first-order and possibly second-order sensitivities of objective and constraints,
- a procedure for approximating the Hessian matrix,
- an efficient QP solver (typically developed separately).

Note that *globalization* strategies such as line search or trust regions are considered out of scope for this paper, because globalization is not always recommended in a real-time setting, since runtime bounds can not be established in general [21], and the initializations are typically close to the exact solution [41]. Possibly extending `acados` with globalization strategies such as

merit functions [59] and filter based line search methods [93] for use in a fully converged setting is subject of future work.

In summary, the SQP algorithm in `acados` looks as follows:

$$\begin{aligned} w^{[i+1]} &= w^{[i]} + \Delta w^{\text{QP}}, & i = 0, 1, \dots, \\ \pi^{[i+1]} &= \pi^{\text{QP}}, & i = 0, 1, \dots, \\ \lambda^{[i+1]} &= \lambda^{\text{QP}}, & i = 0, 1, \dots, \end{aligned}$$

where $w^{[i]} = [x_0^{[i]\top}, u_0^{[i]\top}, \dots, u_{N-1}^{[i]\top}, x_N^{[i]\top}]^\top$ is the primal iterate at SQP iteration i , π^i and λ^i are the dual iterates, readily available from the QP solution. Note that the algebraic variables have been eliminated from the OCP, but numerical approximations of these values are accessible from the numerical integration routine. The step Δw^{QP} is computed by solving the following QP:

$$\begin{aligned} \underset{\substack{\Delta x_0, \dots, \Delta x_N, \\ \Delta u_0, \dots, \Delta u_{N-1}, \\ s_0, \dots, s_N}}{\text{minimize}} \quad & \sum_{k=0}^{N-1} \begin{bmatrix} \Delta x_k \\ \Delta u_k \end{bmatrix}^\top \overbrace{\begin{bmatrix} Q_k & S_k^\top \\ S_k & R_k \end{bmatrix}}^{H_k} \begin{bmatrix} \Delta x_k \\ \Delta u_k \end{bmatrix} + \begin{bmatrix} q_k \\ r_k \end{bmatrix}^\top \begin{bmatrix} \Delta x_k \\ \Delta u_k \end{bmatrix} \\ & + \Delta x_N^\top Q_N \Delta x_N + q_N^\top \Delta x_N \\ & + \sum_{k=0}^N s_k^\top P_k s_k + p_k^\top s_k \end{aligned} \tag{5}$$

subject to

$$\begin{aligned} \Delta x_{k+1} &= A_k \Delta x_k + B_k \Delta u_k + \bar{\phi}_k^x - x_{k+1}, & k = 0, \dots, N-1, \\ -\bar{g}_k &\geq G_k^x \Delta x_k + G_k^u \Delta u_k - J_{s,k} s_k, & k = 0, \dots, N-1, \\ -\bar{g}_N &\geq G_N^x \Delta x_N - J_{s,N} s_N, \\ 0 &\leq s_k & k = 0, \dots, N. \end{aligned}$$

Above, we used the shorthands $\bar{\phi}_k^x$, \bar{g}_k , $k = 0, \dots, N-1$ and \bar{g}_N to denote the function evaluations $\bar{\phi}_k^x(x_k, u_k)$, $g_k(x_k, z_k, u_k)$ and $g_N(x_N)$, respectively. We formulate the NLP in such a way that the slack variables appear directly as tailored optimization variables. We note that some, but not all QP solvers can deal with slack variables directly. For those that do not, slack variables are reformulated as extra input variables.

In the remainder of this section, we discuss each of the ingredients of an efficient SQP solver, described above, separately. Generation of sensitivities using numerical integration will be treated in Section 2.5, Hessian approximation in Section 2.6, structure-exploiting QP solvers in Section 2.7 and real-time considerations in Section 2.8.

2.5 Numerical simulation and sensitivities

An important part of the implementation of direct shooting methods for optimal control consists of reliably and efficiently computing numerical simulation and sensitivity results for the nonlinear system of differential-algebraic equations that represents a dynamic model for our particular system of interest.

Within the family of single-step integration methods one typically distinguishes between explicit and implicit schemes [44]. Well-known examples of explicit integration schemes include explicit Runge-Kutta (RK) formulas such as explicit Euler and the RK method of order 4. Explicit integration schemes are easy to implement since they rely on a direct combination of explicit evaluations of the right-hand side of the system dynamics. Instead, implicit integration schemes result in a nonlinear system of equations that implicitly defines the numerical simulation result. Unlike explicit integration methods, the nonlinear system in implicit integration schemes generally needs to be solved numerically using an iterative procedure such as a Newton-type method. However, implicit formulas are very popular in practice because of their improved numerical stability properties and higher order of accuracy. Especially in case of stiff dynamical systems and implicit or differential-algebraic equations, an implicit integration scheme should often be used [45].

When using these numerical integration schemes within direct multiple shooting, one additionally needs a computationally efficient and reliable way of computing first (and possibly second) order derivatives of the simulation results with respect to the state and control input values:

$$\frac{\partial \phi_k^x(x_k, u_k)}{\partial x_k}, \quad \frac{\partial \phi_k^x(x_k, u_k)}{\partial u_k}, \quad \sum_{i=1}^{n_x} \pi_{k,i} \frac{\partial^2 \phi_{k,i}^x(x_k, u_k)}{\partial^2(x_k, u_k)}, \quad (6)$$

where $\pi_k \in \mathbb{R}^{n_x}$ is called the seed vector, for which the Lagrange multipliers are used to compute the exact Hessian of the Lagrangian. Sensitivity propagation for direct optimal control methods is typically based on a *discretize-then-differentiate* type of approach such as internal numerical differentiation (IND) in [11]. For the class of explicit integration methods, this concept leads to a forward or backward sensitivity propagation based on algorithmic differentiation (AD) techniques [40]. In case of an implicit integration scheme, the IND approach either results in iterative differentiation techniques or a direct computation of sensitivities based on the implicit function theorem [1]. In addition, forward-backward propagation schemes can be derived to compute the symmetric Hessian contributions [78].

Recent work in [77, 76] proposed an algorithmic approach to embed implicit integration schemes with sensitivity analysis in Newton-type optimization for direct optimal control without the need for any iterative procedure, based on the concepts of numerical *condensing* and *expansion* in a lifted Newton-type optimization method [2].

2.6 Convex Hessian Approximation Methods

Gauss-Newton Hessian approximation. In the case of a (nonlinear) least-squares objective in (3), e.g. when tracking a reference, we have

$$l_k(x_k, u_k) = \|r_k(x_k, u_k)\|_2^2, \quad k = 0, \dots, N - 1$$

$$M(x_N) = \|r_N(x_N)\|_2^2,$$

with $r : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_{r_k}}$, $r_N : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_{r_N}}$. Notice that this kind of residual function is a common case in embedded optimization.

The Gauss-Newton Hessian approximation amounts to linearizing "between the norm signs". Since no second-order sensitivities are necessary, the Gauss-Newton Hessian approximation offers a competitive alternative to SQP with exact Hessians, although it converges linearly only. We remark that for a quadratic objective function in (3), the same quadratic objective arises in (5), and no additional computations are needed. For more details on Gauss-Newton methods in the context of NMPC, we refer the reader to [42].

Sequential Convex Quadratic Programming (SCQP). A generalization to using SQP with a Gauss-Newton Hessian approximation is sequential convex quadratic programming [90]. In a sense, a Gauss-Newton SQP algorithm neglects any curvature present in the inequality constraints by linearizing them. In practice however, convex-over-nonlinear objectives and/or constraints arise often, which are of the form $\varphi(c(x, u))$ with a convex function $\varphi(\cdot)$ and a nonlinear function $c(\cdot)$. Examples include ellipsoidal terminal constraints to ensure stability of an NMPC scheme, the friction ellipse in automotive applications, or tunnel-following for robotic manipulators.

In SCQP, we still linearize the inequalities, but bring the convex contributions from the inequality constraints, multiplied with a Lagrange multiplier, into the Hessian approximation. Doing so, the SCQP Hessian is guaranteed to be positive semi-definite. For problems that feature convex-over-nonlinear constraints, this Hessian contribution offers better convergence guarantees than a Gauss-Newton Hessian [90].

Structure-preserving convexification with minimal regularization. In the last two paragraphs, we devised two Hessian approximations which are convex *by construction*. However, when the exact Hessian of the Lagrangian is used, it might be indefinite. When this happens, the optimal direction Δw_{QP} cannot be guaranteed to be a descent direction. Furthermore, many QP solver codes expect a positive (semi-)definite Hessian, even if the second order sufficient conditions for optimality are met. The aim of regularization is to obtain an approximation $\tilde{H} = \text{blkdiag}(\tilde{H}_0, \dots, \tilde{H}_N)$ with each $\tilde{H}_k > 0$. We very briefly discuss three different methods here and compare their convergence in an SQP-type setting in Section 5.

Let $V_k D_k V_k^{-1}$ be the eigenvalue decomposition of H_k , for $k = 0, \dots, N$. Two simple ways of regularizing the Hessian blocks are

$$\tilde{H}_k = \text{project}(H_k, \epsilon) := V_k [\text{maxdiag}(\epsilon, D_k)] V_k^{-1}, \quad (7a)$$

$$\tilde{H}_k = \text{mirror}(H_k, \epsilon) := V_k [\text{maxdiag}(\epsilon, \text{abs}(D_k))] V_k^{-1}, \quad (7b)$$

with ϵ small, $\text{abs}(\cdot)$ defined elementwise and $\text{maxdiag}(\cdot, \cdot)$ selecting the maximum values on the diagonal while not changing the off-diagonal elements.

Another approach is to exploit the optimal control problem structure of (5). One approach of this kind is called *convexification* and has been proposed in [92]. The difference with more naive regularization methods as stated above, is that it first exploits as much convexity as possible through the optimal control structure, before regularizing the remaining negative directions. The complexity of the regularization method is linear in the horizon length, and under some conditions, the SQP iterates converge quadratically to a local solution. An efficient implementation of this new Hessian regularization method is included in `acados`. A numerical example is given in Section 5, which shows the superior convergence behavior of convexification with respect to the mirror and project regularization techniques.

Further convex Hessian approximations Other Hessian approximations exist in the literature, such as BFGS and its limited-memory variant (see [69]). They might be added to `acados` in the future.

2.7 Structure-exploiting embedded QP solvers

There exist different solution strategies for QP (5), which we briefly describe in this section. We note that linear-quadratic optimal control problems can be efficiently solved with the same kind of QP solving strategies presented. As such, `acados`, conceived as a modular software package, can also be used to facilitate solving linear-quadratic QPs, arising in linear-quadratic MPC.

In the following subsections, we give an overview on sparsity exploitation for optimal control structured QPs.

Sparse approach. OCP (5) can be solved directly by using a general-purpose sparse QP solver, e.g., `CVXGEN` [67], `OOQP` [38], both primal-dual interior-point solvers, or `OSQP` [87], a first-order method. First-order methods are mainly based on either the fast gradient method or the alternating direction method of multipliers (ADMM). The strict real-time requirements for solution methods make first-order methods a viable candidate. However, they might suffer from slow convergence rates. For an overview of first-order methods in the context of embedded optimal control, see [27, 56].

Structured approach. OCP (5) is solved by exploiting its multi-stage structure, but dense linear algebra is used. An example is the approach from [86, 80] and in solvers like `FORCES` [24, 96], `HPMPC` [37] and `HPiPM` [34], which are all interior-point solvers.

Condensing approach. An alternative to the previous approaches is the so-called *condensing* approach [13]. By eliminating the state variables by means of the dynamic equality constraints in (5), we obtain a smaller QP with only the control inputs and possibly the initial state as optimization variables. Any general-purpose dense QP solver can then be used to solve the smaller QP, e.g. an active-set solver like qpOASES [28] or an interior-point method like the dense variant in HPIPM. Since qpOASES is able to reuse information (i.e. *warm-starting*) from one problem to the next, it is particularly well-suited for (N)MPC. Condensing is shown to be of quadratic complexity in the horizon length [33].

Partial condensing. A mix between the two previous approaches (i.e. structured, condensing) for solving (5) can be obtained by not eliminating all state variables, but only per blocks of N/N_2 stages (we assume for simplicity that N is an integer multiple of N_2), where N_2 is the ‘new’ horizon length of the partially condensed problem. By this additional degree of freedom, partial condensing enables us to find a better trade-off between horizon length and number of optimization variables, for a given problem. For more details on partial condensing, the reader is referred to [6].

Further condensing strategies. There are further condensing strategies in the literature, which might be added to acados in the future. A method called *complementary condensing* was proposed in [54], where the KKT system is solved in the space of the multipliers corresponding to the equality constraints. This approach is favorable if the problem has more control inputs than states.

Dual-Newton strategy. The dual Newton strategy is an algorithm that is based on dual decomposition tailored to linear-quadratic OCPs in the form of (5), with an open-source implementation in the software qpDUNES [30].

A main advantage of the dual Newton strategy, as in most active set methods, is warm-starting. Contrary to qpOASES, qpDUNES can perform multiple active set changes per iteration. However, a premature termination of the algorithm does not return a meaningful solution as in the case of qpOASES since it is not feasible nor optimal for a neighboring problem. As recently observed in [57], the convergence of qpDUNES can benefit significantly from a partial condensing preprocessing step.

2.8 Real-time iterations

In a real-time control setting, we solve NLP (3) in sequence and under stringent time conditions. Since the environment is anyway changing continuously, it is often sufficient to solve it approximately – it is of no use to have a high-accuracy but past-the-deadline solution.

One such online method is the real-time iteration (RTI) scheme [21]. It solves an inequality-constrained QP in each iteration. The resulting generalized

predictor is better suited for predictions across active set changes, than e.g. a tangential predictor obtained from an interior-point method. For a brief overview, we refer the reader to [22].

In each RTI, one full iteration of an SQP-type scheme is performed, including generation of the sensitivities w.r.t. all variables. We could introduce additional approximations to save on computation time by not updating all sensitivities in each RTI [98]. Such approximations exist on different levels: from only updating the initial state constraint with the current estimate of the state of the system, over updating the right-hand side of the (in)equality constraints, to the full RTI. By interleaving different approximations at different sample times, we obtain a *multi-level* iterations scheme, as introduced by [12].

One important aspect of RTI is initialization. Subsequent RTIs can profit from initialization strategies by warm-starting the algorithm with a (possibly shifted) solution from the previous problem. Recently, an improved and cheap initialization strategy was proposed: Advanced Step Real Time Iteration (AS-RTI) [70]. `acados` was used for its implementation.

3 Algorithm implementations in `acados`

In this section, we focus on the algorithm implementations in `acados`. Since `acados` builds on other software to handle basic linear algebra operations (BLASFEO [35]) and QPs (HPIPM [34]), these will be presented first. Afterwards, a short description of integrators and SQP-type optimization solvers will be given.

3.1 Linear algebra: BLASFEO

At the heart of all embedded optimization routines lies either an implementation of a small set of linear algebra routines (e.g. matrix-matrix multiplication, Cholesky decomposition), or a call to a specialized linear algebra library (e.g. BLAS and LAPACK). Generally BLAS implementations focus on performance for large dense matrices, as used in high-performance computing and data science. Considerably less investigated are BLAS and LAPACK implementations for small dense matrices.

Often, the linear algebra code in embedded optimization packages, for example in the ACADO Code Generation tool or Forces Pro, is code-generated. For very small matrix sizes (e.g. 4×4), this technique outperforms optimized linear algebra libraries. Furthermore, code-generation has the advantage that the code can be kept ‘library-free’. However, for larger matrix sizes (e.g. in the range 10×10 to 100×100 , typical in MPC applications), code-generated linear algebra routines underperform with respect to optimized libraries.

BLASFEO [35] is a linear algebra package that targets computations for small matrices. It offers highly optimized linear algebra routines (e.g. `dgemm`,

`dsyrk`, `dpotrf`), tailored for the matrix sizes typically encountered in embedded optimization. These routines exploit architecture-specific vector instructions for floating point operations (e.g. AVX), and focus on performance for matrices fitting in cache.

Furthermore, BLASFEO defines a packed matrix format (called panel major) which optimizes the cache usage, guaranteeing close to peak performance for matrices of sizes up to a couple hundreds. All high-performance BLASFEO routines use this panel major matrix format, and there is a rich set of auxiliary routines to operate on this matrix format, as well as to convert from/to column- or row-major formats. In this sense, BLASFEO provides a complete linear algebra framework, which can be used to implement many fast optimization algorithms.

Except for trivially small matrices, BLASFEO enables a considerable speedup (up to $10\times$ for some matrix sizes) in the matrix computations, compared to code-generated linear algebra kernels. For all small matrix sizes up to, say, 300×300 , BLASFEO offers a considerable speedup compared to state-of-the-art BLAS implementations, like OpenBLAS, too. The use of BLASFEO is one of the factors why acados performs better than ACADO on medium-scale problems, as we will see in Section 5.

3.2 Quadratic programming: HP-IPM

In the implementation of SQP-type algorithms, QP sub-problems need to be solved efficiently at each iteration. The QP sub-problem solution is typically one of the two most expensive steps in SQP schemes, the other being the simulation and sensitivity computations of dynamical systems.

HP-IPM [34] is a library defining three QP types (dense QP, OCP QP and tree-structured OCP QP, all handling soft constraints via tailored slack variables), and a rich set of routines to create, manage and solve the QPs. All QP solvers are Mehrotra’s type primal-dual interior point methods, and they are implemented using the BLASFEO linear algebra framework.

Furthermore, HP-IPM provides a set of efficiently implemented routines to convert between the different QP types, e.g. condensing routines convert an OCP QP into a dense QP, partial condensing routines convert an OCP QP into another OCP QP with shorter horizon length. Additional expansion routines recover the original QP solution from the (partially) condensed one.

In acados, the QP framework is based on HP-IPM, in the sense that HP-IPM provides both the dense QP and OCP QP definitions, as well as (partial) condensing algorithms to convert them and interior-point methods to solve them. Numerous other QP solvers are then interfaced to acados to alternatively solve the same types of QP problems. At the time of writing, additional interfaces exist to HPMPC, qpDUNES, qpOASES, and OSQP.

3.3 Algorithmic Differentiation

The main difference between software packages for linear quadratic MPC and nonlinear MPC is that NMPC packages need a modeling language for the nonlinear functions. We support `CasADi`, a graph-based source transformation algorithmic differentiation tool [4]. In our workflow, a user would typically specify the dynamic continuous-time models and nonlinear constraint functions with `CasADi` in a high-level language such as `MATLAB` or `Python`. In combination with `CasADi`'s code generator we obtain fast, embeddable code, see Section 4.3.

3.4 Numerical simulation

`acados` features different kinds of numerical simulation routines. There are implementations of explicit and implicit Runge-Kutta integrators available, both of which support the optional propagation of first-order forward and adjoint sensitivities, as well as second-order sensitivities. The explicit integrators can be used with explicit ODE models and supports different Butcher tableaus, including Euler's method and RK4. Moreover, the implicit integrators can be used with an index-1 differential-algebraic equation (DAE) or implicit ODE model and use the Gauss-Legendre Butcher tableaus. A novel implementation for lifted collocation integrators [77] has been made part of `acados`, as well as a recently proposed structure-exploiting IRK algorithm, the so-called GNSF-IRK scheme [32], discussed next.

The concept of GNSF-IRK is to rigorously exploit the linear dependencies within the dynamic system. It extends the ideas of the linear input and linear output subsystems that have been implemented within the `ACADO Code Generation` tool [75] and uses a more flexible structured dynamic system formulation that can also handle index-1 DAEs. A main challenge for structure-exploiting integrators is to appropriately reformulate the dynamic system of interest into the desired structured form. `acados` features an automatic transcription method for the GNSF structure [32], implemented as a `MATLAB` function for `CasADi` models.

A last important feature of `acados` is that integrators can vary from stage to stage, with e.g. different state and control dimensions, different integration step length, or different integration schemes.

We note that all integrator modules, except for the ERK integrator, are based on hardware tailored linear algebra routines in `BLASFEO` to speed up the LU factorizations and the corresponding triangular system solutions, as discussed in [36].

3.5 SQP-type methods

For nonlinear programming, `acados` offers different SQP-like methods. A full-step SQP method is available, with different algorithmic options. As Hessian

approximations, we have Gauss-Newton Hessians, SCQP, and exact Hessians with regularization/convexification as discussed in Section 2.8, and we allow for user-defined Hessian approximations. For use in an online setting, e.g. in NMPC, a specialized RTI routine is available.

Both the SCQP algorithm and the convexification method of Section 2.6 are novel features, to the authors’ knowledge, not present in any other NMPC software packages.

4 The `acados` software package

`acados` implements some of the optimization methods mentioned in the previous sections. `acados` is meant to be user-friendly at a high level, and efficient at a low level. In order to balance these properties, we developed a core library written in C which exposes functionality to the `Python` and `MATLAB` interfaces. In this section, we first discuss the functionality of this inner core module, we then describe internal and external interfaces that are crucial for usability.

4.1 The `acados` core library

The embedded optimization algorithms discussed in Section 2 are implemented in `acados` in a modular fashion. For example, there is a clear interface between an NLP solver and an integrator. The integrator expects a linearization point $w^{[2]}$ and returns the end state of a simulated trajectory, and optionally first- and second-order sensitivities:

$$\text{NLP solver} \begin{array}{c} \xrightarrow{\text{lin. point}} \\ \xleftarrow{\text{sim.,sens.}} \end{array} \text{integrator}$$

Similar diagrams can be drawn for all other algorithmic components, including (partial) condensing, QP solvers, function evaluations etc. Each of these algorithmic components are modeled within `acados` as separate *modules*. Some modules can be used as standalone modules, or in combination with others. For instance, depending on the choice of algorithm, an NLP solver will make use of some or all of the other modules. In Table 2, we see an overview of all modules currently present in `acados`, together with the implemented algorithmic variants.

It is an important design choice that all modules are identical in their signature. That way, all modules look similar to the users of `acados`. For developers, it should be straightforward to extend `acados` with another module. The signature is as follows (in C syntax):

```
int <solver>(void *config,
            void *dims,
            <module>_in *in,
            <module>_out *out,
            void *opts,
```

Table 2 OVERVIEW OF THE SOFTWARE MODULES PRESENT IN ACADOS.

Module	Variants
OCP QP	HPIPM qpDUNES HPMPC OSQP
Dense QP	HPIPM qpOASES
Condensing	Full condensing (HPIPM) Partial condensing (HPIPM)
Simulation	ERK IRK GNSF-IRK lifted IRK
OCP NLP	Gauss-Newton SQP Gauss-Newton SCQP Exact-Hessian SQP RTI
Regularization	Projection Mirroring Convexification
Nonlinear function	CasADi generated functions C-code functions

```
void *mem,
void *work);
```

Here, `<module>` stands for the name of the module at hand, for example `ocp_qp` for QP problems with optimal control structure or `sim` for integration problems, and `<solver>` is a placeholder for a function implementing the specific solver for problems corresponding to this module, e.g., `ocp_qp_hpipm` (interface to HPIPM solver) or `sim_erk` (explicit Runge-Kutta method), etc. Each module returns an `int` which denotes a solver-specific error status – zero means successful completion by convention. All input arguments are pointers. Each of the arguments comes with a set of helper functions, called `...calculate_size`, computing the size (in bytes) of the `struct` pointed to, as well as a set of functions, called `...assign`, to initialize a block of memory.

Some modules comprise other modules. For example, an SQP solver for optimal control problems might need an integrator, which is on its own a proper `acados` module. In this context, we call the integrator a *submodule*. Each of the arguments above, `dims`, `in`, etc., have fields corresponding to submodules. As an example, the relation between an NLP solver and its submodules is depicted in Figure 1. We remark that the calculation of the memory size of a module with submodules is done recursively, i.e., calling the `calculate_size` function on the top module returns the required memory size of the top module and all of its submodules, and submodules of submodules, etc. This allows users to allocate all the memory outside of `acados`, by design.

The *core* library of `acados` contains mostly what has been described in this section: a collection of modules, each with corresponding data types and

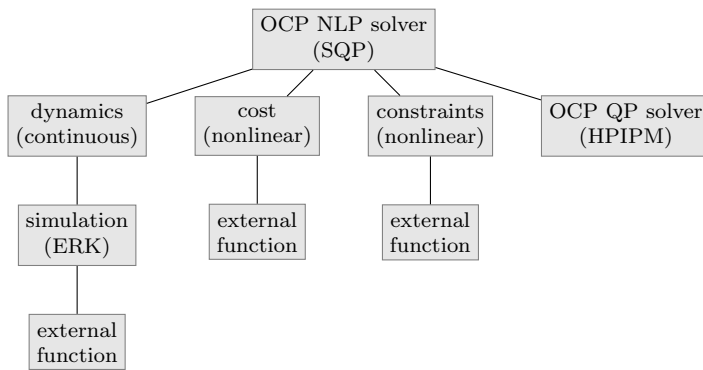


Fig. 1 Example of the relation between modules and submodules in acados for a specific case of a possible SQP algorithm.

variants of solvers, as well as helper functions for memory management. Using the core library directly can be cumbersome and error-prone, as many details need to be taken into account: it is designed to be efficient and flexible. To cater to the specific needs of the end user, we offer different interfaces to the core of acados, which are described next.

4.2 The C interface

The C interface is responsible for encapsulating the low-level constructs of the acados core.

Choosing Solvers. When working with the core library, all functions are specific to one variant of a module: when solving a QP with, say, qpOASES, the code will refer to structs like `dense_qp_qpoases_memory`, `dense_qp_qpoases_opts`, etc. We provide an abstraction layer to facilitate switching solvers easily. To this end, for each module we define a ‘plan’. A plan is a struct that contains a number of fields representing the choice of a particular combination of solvers. For example, the plan for an SQP-type method with Gauss-Newton Hessian approximation, for a problem discretized with an ERK integrator using HPIP as an underlying QP solver, reads as

```

ocp_nlp_solver_plan plan = {
    {PARTIAL_CONDENSING_HPIP},
    {ERK, ERK, ERK, ...},
    SQP_GN,
    {LINEAR_LS, ...},
    {CONTINUOUS_MODEL, ...},
    {BGH, ...},
};
  
```

Here, the arrays should be of the correct length (omitted for brevity, with a slight abuse of notation). As a general rule, solvers that make use of other modules should include them in their plan.

Passing options. For each module, we manipulate a specific `options` struct using functions that take a textual representation of the option via a string. The string encodes both the module that the option belongs to, as well as the name of the option.

Memory management. Allocating memory ‘manually’ as described above can quickly become cumbersome. For this reason, we make available a few routines that automate that process. To this end, in the C interface each module from the core library is mirrored by an additional function with signature

```
<module>_solve(<module>_solver *solver,
               <module>_in *in,
               <module>_out *out);
```

The argument `solver` is a pointer to a C structure that encapsulates the data needed other than input and output. By doing so, we reduce the amount of boilerplate code.

Convenience routines. The C interface additionally offers helper routines, so-called ‘setters’ and ‘getters’, that wrap the handling of the low-level structs of the `acados` core.

4.3 High-level interfaces

Often, software for NMPC is coded in scripting languages. Therefore, we offer interfaces to two popular languages for scientific computing: Python and MATLAB, where the interface to MATLAB is largely compatible with its free and open-source alternative Octave. As such, we created a small domain-specific language within each of these frameworks. We build on top of code from the C interface of `acados`.

In order to formulate the OCP (3) through the `acados` modules (`cost`, `constraints` and `dynamics`), the main challenge is to pass the generally nonlinear functions and their derivatives to these modules. The Python and MATLAB interfaces of `acados` use `CasADi` as a modeling language, i.e. to formulate all generally nonlinear parts of the OCP. The `acados` high level interfaces are able to use `CasADi`’s code generation and algorithmic differentiation to generate the C functions needed for each `acados` module. A major benefit of using `CasADi` as a modeling language is that the solution behavior of `acados` can be easily compared with the solutions coming from the numerous optimization tools interfaced with `CasADi`.

Once the OCP to be solved is described through the domain-specific language implemented by the high-level interfaces, a human readable self-contained C project that makes use of templated code can be generated. The generated project contains all the C code necessary for function and derivative evaluations generated through `CasADi` and the C code necessary to set up the NLP solver using the `acados` C interface. Moreover, a MATLAB S-Function and a build system for its compilation is generated. Note that this kind of code

generation is inherently different from the one in ACADO, since the templated code uses only the functions exposed by the C interface of `acados`. In contrast to this, ACADO generated solvers are standalone C projects that are extremely problem specific and do not rely on a common library.

With the workflow described above, it is possible to obtain a self-contained, high-performance solver that can be easily deployed on embedded hardware starting from a description of the OCP in a high-level language.

We remark that model equations and other nonlinear functions are called from `acados` in a completely language-agnostic way: `acados` is at no point aware of which modeling tool is being used. One benefit is that this facilitates self-written models (in C/C++), which are also completely compatible with `acados`. However, this is more involved, since in the case of `CasADi` functions, memory allocation and matrix format conversions are taken into account automatically by the `CasADi` functions wrapper in `acados`.

5 Numerical Results

This section consists of a few numerical experiments with `acados` and comparisons to other embedded optimization software packages. We discuss performance on the nonlinear chain-of-masses problem, we present one open-loop example with different Hessian approximations, and show one closed-loop engine control experiment on an embedded platform.

5.1 Case Study 1: Chain of Masses

As a benchmarking problem, we take the chain-of-masses problem as presented in [95]. The control objective is to stabilize a chain of masses with nonlinear interaction between them. For a full description of the system, we refer to the appendix. The system is useful as a benchmark in the sense that the problem is simple enough to understand intuitively, yet complicated enough to get non-trivial results from a range of different solvers. Also, by increasing the number of masses, one could compare behavior for different numbers of states easily, without changing much code.

Closed-loop experiments

In closed-loop, an MPC controller repeatedly (approximately) solves OCP (10). The first control u_0 is passed to the dynamic system under control and a new initial state \bar{x}_0 is obtained. Here, we simulate the system by using a more accurate integrator than the one in OCP (10), namely the Dormand-Prince method, as implemented in the MATLAB routine `ode45`.

We introduce one disturbance into the closed-loop system, similar as in [95]: in the beginning of the simulation, we start from a horizontal configuration of the chain of masses. Around the midpoint of the simulation, we override

the closed loop control with a constant $u_d = [-1, 1, 1]^\top$. After one second of simulation time, the controller takes over again.

We compare the following solvers with each other for this particular closed-loop setup:

- IPOPT [94]. As a solver *not* targeting embedded devices specifically, we use it as a baseline to compare against.
- FalcoPT [88]. A projected gradient descent method tailored for NMPC.
- VIATOC [50]. A gradient projection method for MPC that only allows linear inequality constraints.
- ACADO Code Generation tool [49]. Generates SQP-based solvers.
- GRAMPC [26]. An embedded Augmented Lagrangian-based solver.
- acados. Framework presented in the current paper.

The tuning parameters for the different solvers are listed in Table 3.

Table 3 SOLVER OPTIONS FOR THE DIFFERENT SOLVERS IN CASE STUDY 1.

Solver	Solver options
IPOPT	Called through CasADi, default parameters
FalcoPT	Tolerance (eps): 0.1, maximum number of iterations (maxIt): 100
VIATOC	Maximum number of iterations: 20
ACADO	RTI solver, Full condensing, QP solver qpOASES
GRAMPC	Parameters chosen as in [26]: max. number of augm. Lagrange iterations: 5
acados	SQP_RTI solver, QP solver HPiPM, partial condensing horizon of 5

In order to compare the quality of the closed-loop solutions, we use the notion of distance-from-reference (DR), which is an approximation of the integrated cost along closed-loop trajectories:

$$\text{DR}_{(\cdot),n} = \sum_{i=0}^n \begin{bmatrix} x_{(\cdot),i} - x_{\text{ref}} \\ u_{(\cdot),i} - u_{\text{ref}} \end{bmatrix}^\top \begin{bmatrix} Q & 0 \\ 0 & R \end{bmatrix} \begin{bmatrix} x_{(\cdot),i} - x_{\text{ref}} \\ u_{(\cdot),i} - u_{\text{ref}} \end{bmatrix}.$$

To compare the different solvers, we plot the relative cumulative suboptimality (RCSO), relative to a fully converged solution, in this case, the IPOPT solution, which reads as

$$\text{RCSO}_{(\cdot),n} = \left| \frac{\text{DR}_{(\cdot),n} - \text{DR}_{\text{ipoppt},n}}{\text{DR}_{\text{ipoppt},n}} \right|,$$

where $n = 0, \dots, 300$ denotes the time step in our simulation. We show a comparison in Table 4. The results for ACADO and acados are *exactly* the same, as they implement the same real-time algorithm, with both being very close to the reference solution from IPOPT. The solvers GRAMPC, VIATOC and FalcoPT, being based on first-order methods, are further away from the IPOPT solution. These findings are consistent with previously published work by other authors, see [26].

Table 4 Relative suboptimality at the end of the simulation of the hanging chain with $M = 5$ and $N = 40$. First-order methods VIATOC, GRAMPC and FalcoOPT were tuned to perform similarly. The algorithms chosen in ACADO and acados are identical, hence the results are identical.

Solver name	RCSO
IPOPT	0.00e+00
FalcoOPT	3.170e-01
VIATOC	4.74e-03
ACADO	1.01e-04
GRAMPC	7.17e-02
acados	1.01e-04

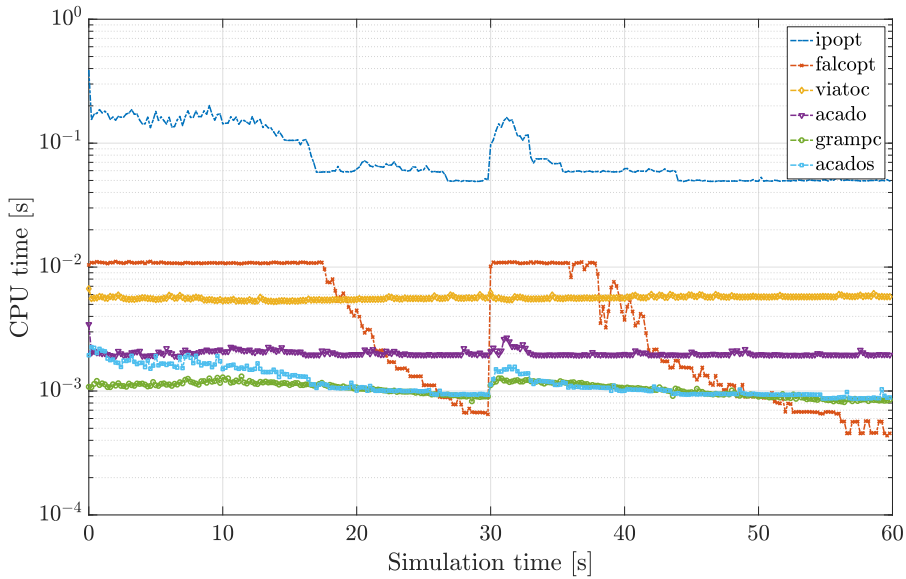


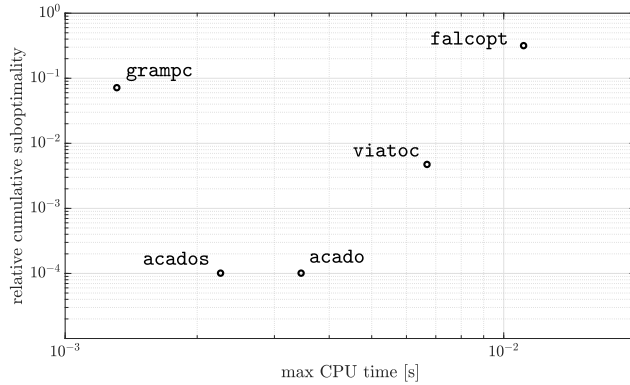
Fig. 2 Computational time for each iteration of the closed loop simulation, averaged over 10 runs.

We have a look at the computational performance along the closed-loop trajectories in Figure 2. GRAMPC, ACADO, VIATOC and acados produce consistent timings throughout the entire experiment, even when the disturbance occurs. This is a beneficial property for embedded solvers, as they often have a fixed time deadline, being part of a larger control application. GRAMPC and acados produce solutions at almost the same speed, both approximately a factor 2 faster than ACADO which is in turn a factor 2-3 faster than VIATOC. Near the equilibrium, FalcoOPT takes the shortest computation time, as it needs to perform only a few gradient steps per iteration. IPOPT is included as a baseline for comparison to non-embedded solvers. The timings are summarized in Table 5.

Of course, any solver can trade off sub-optimality for computation time. To get the full picture, we plot both measures against each other in Figure 3:

Table 5 COMPUTATION TIMES FOR THE CLOSED-LOOP EXPERIMENTS ON A CHAIN OF MASSES (CF. FIGURE 2)

comp. time per iteration (ms)	median	minimum	maximum
IPOPT	59.84	49.06	384.90
FalcoPT	4.36	0.44	11.10
VIATOC	5.63	5.27	6.68
ACADO	1.97	1.90	3.45
GRAMPC	1.06	0.81	1.31
acados	1.05	0.87	2.23

**Fig. 3** Trade-off between sub-optimality (Table 4) and computation time (Figure 2). We see that acados and GRAMPC lie on the Pareto-optimal front.

we look at relative cumulative sub-optimality over the entire length of the experiment, versus worst-case computation times. By this comparison, we see that acados and GRAMPC are on the Pareto-optimal front: although acados is a factor 1000 less suboptimal than GRAMPC, the computational cost is higher. By the median computation times, acados is faster (see Table 5).

5.2 Case study 2: Hessian regularization

In Section 2.8, we briefly mentioned the impact of Hessian regularization on SQP methods. In this case study, we compare the convergence of exact-Hessian based SQP with three different Hessian regularizations, on a simple control problem, namely a cart-pole swingup, which is described in the appendix. Note that in this case study, as opposed to the others, we do not perform closed-loop experiments but solve one optimal control problem up until convergence.

5.2.1 Exact-Hessian based SQP

We solve OCP (12) with SQP, where we use the exact Hessian of the Lagrangian. In the notation of (5):

Table 6 EXACT-HESSIAN BASED SQP: COMPUTATION TIMES

regularization	convexification	project	mirror
avg iteration (ms)	2.540	2.303	2.264
total time (ms)	66.034	103.65	174.32

$$H_k = \begin{bmatrix} Q & 0 \\ 0 & R \end{bmatrix} + \sum_{i=0}^{n_x} \pi_{k,i} \nabla_{(x,u)}^2 \phi_i^x(x_k, u_k), \quad k = 0, \dots, N - 1$$

$$H_N = Q,$$

where $\pi_{k,i}$ are the Lagrange multipliers associated with the dynamic equality constraints.

In some cases, the non-convexity of the dynamic equations gives rise to an indefinite Hessian matrix. We apply the `project(·)` and `mirror(·)` regularizations, as well as the convexification method previously mentioned in Section 2.6. All are implemented as modules in the `acados` framework.

We compare the convergence of the SQP iterates obtained using the three different regularization methods. For each SQP variant, we start the SQP iterations from an initialization point with zeros for all states except a linearly decreasing initialization for the angle, from π to zero. The result can be seen in Figure 4. The structure-exploiting convexification converges almost twice as fast as the projection regularization, and is in turn much faster than the mirroring regularization. Intuitively, this makes sense, as mirroring is ‘blocking’ directions associated with large negative eigenvalues, by introducing large positive eigenvalues in those directions. This prevents the solver from taking larger steps². In turn, the structure-exploiting regularization is faster than merely projecting the eigenvalues on the positive definite cone, because it is redistributing convexity among all stages, and thus needs less regularization overall.

It must be said that the convexification method is quite a bit more involved than the other two regularization schemes. However, by using the optimized linear algebra routines of BLASFEO, we implemented the convexification method such that it is only slightly more expensive per iteration than the basic regularization methods, see Table 6, but much less computationally expensive overall. Thus, the Hessian convexification method allows us to perform exact-Hessian based NMPC online, with better performance than state-of-the-art methods.

² This is also the approach followed by the algorithms obtained with the `ACADO Code Generation` tool.

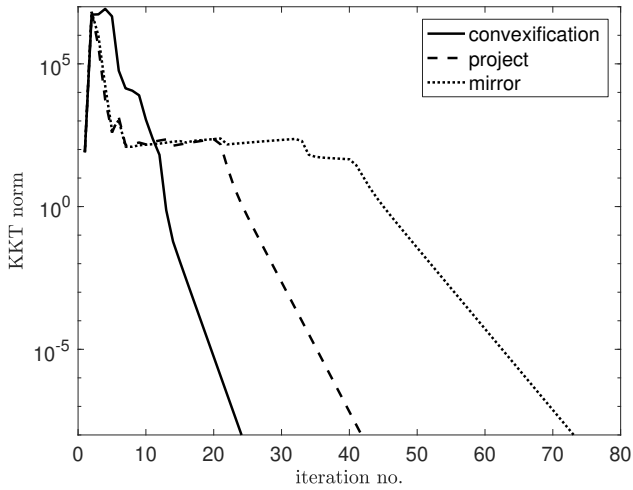


Fig. 4 Convergence comparison of exact-Hessian based SQP with three different regularization strategies.

5.3 Case study 3: Hardware-in-the-loop experiments for an engine control application

As a last case study, we discuss the performance of `acados` on an embedded platform, namely the dSPACE MicroAutoboxII [25]. It is an industrial computing platform that is used in the car industry. It features a 900 MHz PowerPC processor (IBM PPC 750GL) with 16MB of main memory. The control application that we focus on is engine control, with the engine model as presented in [3], which we will briefly reproduce in the appendix.

The control objective is to track a boost pressure signal, where the boost pressure is given by $y_p(x) := \Pi_{c,lp} \cdot \Pi_{c,hp}$ (see appendix). To this end, we solve an OCP arising from a multiple shooting formulation with the Gauss-Legendre method of order 6 with sampling time 0.05 s and $N = 20$ shooting intervals. The DAE simulation functions are denoted by ϕ . Let $r(x, u) = [y_p(x); x; u]$ and $r_N(x) = [y_p(x); x]$. The OCP then reads as

$$\begin{aligned}
 & \underset{\substack{x_0, \dots, x_N, \\ z_0, \dots, z_{N-1}, \\ u_0, \dots, u_{N-1}}}{\text{minimize}} && \sum_{k=0}^{N-1} \|r(x_k, u_k) - y_{r,k}\|_W^2 + \|r_N(x_N) - y_{r,N}\|_{W_N}^2 \\
 & \text{subject to} && x_0 = \bar{x}_0, \\
 & && \begin{bmatrix} x_{k+1} \\ z_k \end{bmatrix} = \phi(x_k, u_k), \quad k = 0, \dots, N-1, \\
 & && 0 \leq u_k \leq 100, \quad k = 0, \dots, N-1, \\
 & && 0.5 \leq \Pi_{c,lp,k} \leq 1.757, \quad k = 1, \dots, N, \\
 & && 0.5 \leq \Pi_{c,hp,k} \leq 2.125, \quad k = 1, \dots, N.
 \end{aligned} \tag{8}$$




Fig. 5 Closed-loop simulation of the engine control task with steps in the reference boost pressure. Simulations are carried out on the dSPACE MicroAutoboxII platform at a clock speed of 900 MHz.

Constraints on $\Pi_{c,lp}$, $\Pi_{t,hp}$ are included to prevent damage to the compressor. The exact values of the weight matrices and the reference vectors can be found in the appendix.

We repeatedly solve OCP (8) approximately by performing real time iterations. As an underlying QP solver, we use `HP1PM`. When ran in closed loop on the dSPACE MicroAutoboxII, the results can be seen in Figure 5. Control bounds and state bounds become active at some point in the simulation, for the high-pressure stage. The reference is tracked closely and without oscillations, which have been observed when linear-quadratic MPC is used [3]. As for the computation times, it is interesting to note that there are spikes everywhere where a jump occurs or a constraint becomes (in)active. The computation times close to the solution (i.e. at the beginning of the simulation) drop to almost zero. Overall, the maximum computation time remains under 10 ms, which is 5x faster than the sampling time of the system (50 ms).

6 Conclusion and outlook

In this article, we presented `acados`, a new software package for embedded optimization. It is free and open-source software that facilitates rapid testing and deployment of (N)MPC algorithms on embedded hardware platforms. For

ease of use, we offer interfaces with higher-level languages such as MATLAB and Python.

Among many features that state-of-the-art NMPC algorithms require, a couple of new features that are not present in any other software package is the convexification procedure of Section 2.6, allowing the use of exact-Hessian based SQP methods in real-time, and the SCQP Hessian approximation. Additionally, the structure exploiting GNSF-IRK integrator has the potential to speed up the simulation and sensitivity propagation tasks within an NMPC scheme. Furthermore, `acados` features partial condensing, different state and control dimensions per multiple shooting stage, the use of BLASFEO as a linear algebra backend and facilities for using CasADi as a modeling language.

The software is shown to be embeddable, by numerical experiments on the dSPACE MicroAutoboxII industrial computer, resulting in computation times in the millisecond range for a non-trivial NMPC problem. Furthermore, it is shown to be fast, by comparison to other embedded optimization packages. `acados` is already being used by others in various control and estimation applications, such as reluctance synchronous machines [97], real-time estimation of muscle forces [7], wind turbine fatigue [64], quadrotors [14], autonomous car racing [55], etc.

`acados` is an ongoing endeavor. Future work includes extending interoperability with Simulink for easier deployment on embedded systems, and adding features for nonlinear interior-point methods, as well as other SQP-based methods like multi-level iterations.

References

1. Albersmeyer, J.: Adjoint-based algorithms and numerical methods for sensitivity generation and optimization of large scale dynamic systems. Ph.D. thesis, University of Heidelberg (2010)
2. Albersmeyer, J., Diehl, M.: The lifted Newton method and its application in optimization. *SIAM Journal on Optimization* 20(3), 1655–1684 (2010)
3. Albin, T., Ritter, D., Liberda, N., Quirynen, R., Diehl, M.: In-vehicle realization of nonlinear MPC for gasoline two-stage turbocharging airpath control. *IEEE Transactions on Control Systems Technology* pp. 1–13 (2017)
4. Andersson, J.A.E., Gillis, J., Horn, G., Rawlings, J.B., Diehl, M.: CasADi – a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation* 11(1), 1–36 (2019)
5. Andersson, J.A.E., Rawlings, J.B.: Sensitivity analysis for nonlinear programming in casadi. In: *Proceedings of the IFAC Conference on Nonlinear Model Predictive Control (NMPC)* (2018)
6. Axehill, D.: Controlling the level of sparsity in MPC. *Systems & Control Letters* 76, 1–7 (2015)
7. Bailly, F., Ceglia, A., Michaud, B., Rouleau, D.M., Begon, M.: Real-time and dynamically consistent estimation of muscle forces using a moving horizon emg-marker tracking algorithm—application to upper limb biomechanics. *Frontiers in Bioengineering and Biotechnology* 9, 112 (2021)
8. Bemporad, A.: *Hybrid Toolbox for Matlab* (2003)
9. Bemporad, A., Borrelli, F., Morari, M.: The explicit solution of constrained LP-based receding horizon control. In: *Proceedings of the IEEE Conference on Decision and Control (CDC)*. Sydney, Australia (1999)

10. Bock, H.: Randwertproblemmethoden zur Parameteridentifizierung in Systemen nicht-linearer Differentialgleichungen, *Bonner Mathematische Schriften*, vol. 183. Universität Bonn, Bonn (1987)
11. Bock, H.G.: Recent advances in parameter identification techniques for ODE. In: Numerical Treatment of Inverse Problems in Differential and Integral Equations, pp. 95–121. Birkhäuser (1983)
12. Bock, H.G., Diehl, M., Kühn, P., Kostina, E., Schlöder, J.P., Wirsching, L.: Numerical methods for efficient and fast nonlinear model predictive control. In: Proceedings of "Int. Workshop on assessment and future directions of Nonlinear Model Predictive Control". Springer (2005)
13. Bock, H.G., Plitt, K.J.: A multiple shooting algorithm for direct solution of optimal control problems. In: Proceedings of the IFAC World Congress, pp. 242–247. Pergamon Press (1984)
14. Carlos, B.B., Sartor, T., Zanelli, A., Frison, G., Burgard, W., Diehl, M., Oriolo, G.: An efficient real-time nmpc for quadrotor position control under communication time-delay. In: 2020 16th International Conference on Control, Automation, Robotics and Vision (ICARCV), pp. 982–989 (2020)
15. Chen, Y., Bruschetta, M., Picotti, E., Beghi, A.: Matmpc - a matlab based toolbox for real-time nonlinear model predictive control. arXiv preprint (2018)
16. Cheshmi, K., Kaufman, D.M., Kamil, S., Dehnavi, M.M.: NASOQ: Numerically accurate sparsity-oriented qp solver. In: ACM Transactions on Graphics, vol. 39
17. Chiang, N.Y., Hang, R., Zavala, V.M.: An augmented lagrangian filter method for real-time embedded optimization. *IEEE Transactions on Automatic Control* 62(12), 6110–6121 (2017)
18. Cimini, G., Bemporad, A.: Exact complexity certification of active-set methods for quadratic programming. *IEEE Transactions on Automatic Control* 62(12), 6094–6109 (2017)
19. Deng, H., Ohtsuka, T.: A highly parallelizable newton-type method for nonlinear model predictive control. In: Proceedings of the IFAC Conference on Nonlinear Model Predictive Control (NMPC) (2018)
20. Di Cairano, S., Brand, M., Bortoff, S.A.: Projection-free parallel quadratic programming for linear model predictive control. *International Journal of Control* 86(8), 1367–1385 (2013)
21. Diehl, M., Bock, H.G., Schlöder, J.P., Findeisen, R., Nagy, Z., Allgöwer, F.: Real-time optimization and nonlinear model predictive control of processes governed by differential-algebraic equations. *Journal of Process Control* 12(4), 577–585 (2002)
22. Diehl, M., Ferreau, H.J., Haverbeke, N.: Efficient numerical methods for nonlinear MPC and moving horizon estimation. In: L. Magni, M. Raimondo, F. Allgöwer (eds.) Nonlinear model predictive control, *Lecture Notes in Control and Information Sciences*, vol. 384, pp. 391–417. Springer (2009)
23. Domahidi, A., Chu, E., Boyd, S.: ECOS: An SOCP solver for embedded systems. In: Proceedings of the European Control Conference (ECC), pp. 3071–3076. IEEE (2013)
24. Domahidi, A., Zraggen, A., Zeilinger, M.N., Morari, M., Jones, C.N.: Efficient interior point methods for multistage problems arising in receding horizon control. In: Proceedings of the IEEE Conference on Decision and Control (CDC), pp. 668–674. Maui, HI, USA (2012)
25. dSPACE: Homepage. <http://www.dspace.com> (2006)
26. Englert, T., Völz, A., Mesmer, F., Rhein, S., Graichen, K.: A software framework for embedded nonlinear model predictive control using a gradient-based augmented lagrangian approach (GRAMPC). *Optimization and Engineering* 20(3), 769–809 (2019)
27. Ferreau, H.J., Almer, S., Verschueren, R., Diehl, M., Frick, D., Domahidi, A., Jerez, J.L., Stathopoulos, G., Jones, C.: Embedded optimization methods for industrial automatic control. In: Proceedings of the IFAC World Congress (2017)
28. Ferreau, H.J., Kirches, C., Potschka, A., Bock, H.G., Diehl, M.: qpOASES: a parametric active-set algorithm for quadratic programming. *Mathematical Programming Computation* 6(4), 327–363 (2014)
29. Franke, R., Arnold, E.: Computer Intensive Methods in Control and Signal Processing, chap. Applying new numerical algorithms to the solution of discrete-time optimal control problems, pp. 105–117. Springer (1997)

30. Frasch, J.V., Sager, S., Diehl, M.: A parallel quadratic programming method for dynamic optimization problems. *Mathematical Programming Computations* 7(3), 289–329 (2015)
31. Frey, J., Cairano, S.D., Quirynen, R.: Active-Set based Inexact Interior Point QP Solver for Model Predictive Control. In: *Proceedings of the IFAC World Congress* (2020)
32. Frey, J., Quirynen, R., Kouzoupis, D., Frison, G., Geisler, J., Schild, A., Diehl, M.: Detecting and exploiting Generalized Nonlinear Static Feedback structures in DAE systems for MPC. In: *Proceedings of the European Control Conference (ECC)* (2019)
33. Frison, G.: Algorithms and methods for high-performance model predictive control. Ph.D. thesis, Technical University of Denmark (DTU) (2015)
34. Frison, G., Diehl, M.: HPIPM: a high-performance quadratic programming framework for model predictive control. In: *Proceedings of the IFAC World Congress*. Berlin, Germany (2020)
35. Frison, G., Kouzoupis, D., Sartor, T., Zanelli, A., Diehl, M.: BLASFEO: Basic linear algebra subroutines for embedded optimization. *ACM Transactions on Mathematical Software (TOMS)* 44(4), 42:1–42:30 (2018)
36. Frison, G., Quirynen, R., Zanelli, A., Diehl, M., Jørgensen, J.B.: Hardware tailored linear algebra for implicit integrators in embedded NMPC. In: *Proceedings of the IFAC World Congress* (2017)
37. Frison, G., Sorensen, H.B., Dammann, B., Jørgensen, J.B.: High-performance small-scale solvers for linear model predictive control. In: *Proceedings of the European Control Conference (ECC)*, pp. 128–133 (2014)
38. Gertz, E.M., Wright, S.J.: Object-oriented software for quadratic programming. *ACM Transactions on Mathematical Software* 29(1), 58–81 (2003)
39. Gifftaler, M., Neunert, M., Stäuble, M., Buchli, J.: The Control Toolbox - an open-source C++ library for robotics, optimal and model predictive control. In: *IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAN)* (2018)
40. Griewank, A.: *Evaluating Derivatives, Principles and Techniques of Algorithmic Differentiation*. No. 19 in *Frontiers in Appl. Math.* SIAM, Philadelphia (2000)
41. Gros, S., Srinivasan, B., Bonvin, D.: Robust predictive control based on neighboring extremals. *Journal of Process Control* 16, 243–253 (2006)
42. Gros, S., Zanon, M., Quirynen, R., Bemporad, A., Diehl, M.: From linear to nonlinear MPC: bridging the gap via the real-time iteration. *International Journal of Control* (2016)
43. Grüne, L., Pannek, J.: *Nonlinear Model Predictive Control*, second edition edn. Springer-Verlag (2017)
44. Hairer, E., Nørsett, S., Wanner, G.: *Solving Ordinary Differential Equations I*, 2nd edn. Springer Series in Computational Mathematics. Springer, Berlin (1993)
45. Hairer, E., Wanner, G.: *Solving Ordinary Differential Equations II – Stiff and Differential-Algebraic Problems*, 2nd edn. Springer, Berlin Heidelberg (1991)
46. Hehn, M., D’Andrea, R.: A flying inverted pendulum. In: *IEEE International Conference on Robotics and Automation*, pp. 763–770
47. Herceg, M., Kvasnica, M., Jones, C., Morari, M.: Multi-Parametric Toolbox 3.0. In: *Proc. of the European Control Conference*, pp. 502–510. Zürich, Switzerland (2013). <http://control.ee.ethz.ch/mpt>
48. Hermans, B., Themelis, A., Patrinos, P.: Qpalm: A proximal augmented lagrangian method for nonconvex quadratic programs
49. Houska, B., Ferreau, H.J., Diehl, M.: An auto-generated real-time iteration algorithm for nonlinear MPC in the microsecond range. *Automatica* 47(10), 2279–2285 (2011)
50. Kalmari, J., Backman, J., Visala, A.: A toolkit for nonlinear model predictive control using gradient projection and code generation. *Control Engineering Practice* 39, 56–66 (2015)
51. Käpernick, B., Graichen, K.: The gradient based nonlinear model predictive control software GRAMPC. In: *Proceedings of the European Control Conference (ECC)* (2014)
52. Katliar, M.: *Optimal control of motion simulators*. Ph.D. thesis, Albert-Ludwigs-Universität Freiburg (2020)
53. Khusainov, B., Kerrigan, E.C., Suardi, A., Constantinides, G.A.: Nonlinear predictive control on a heterogeneous computing platform. In: *Proceedings of the IFAC World Congress* (2017)

54. Kirches, C., Bock, H.G., Schlöder, J.P., Sager, S.: Complementary condensing for the direct multiple shooting method. In: *Modeling, Simulation and Optimization of Complex Processes*, pp. 195–206. Springer Berlin Heidelberg (2012)
55. Kloeser, D., Schoels, T., Sartor, T., Zanelli, A., Frison, G., Diehl, M.: NMPC for racing using a singularity-free path-parametric model with obstacle avoidance. In: *Proceedings of the IFAC World Congress* (2020)
56. Kouzoupis, D., Ferreau, H.J., Peyrl, H., Diehl, M.: First-order methods in embedded nonlinear model predictive control. In: *Proceedings of the European Control Conference (ECC)*, pp. 2617–2622 (2015)
57. Kouzoupis, D., Quirynen, R., Frasch, J.V., Diehl, M.: Block condensing for fast nonlinear MPC with the dual Newton strategy. In: *Proceedings of the IFAC Conference on Nonlinear Model Predictive Control (NMPC)*, vol. 48, pp. 26–31 (2015)
58. Kvamme, S.: DuQuad Webpage. <http://sverrka.github.io/duquad/> (2014)
59. Leineweber, D.B.: Efficient reduced SQP methods for the optimization of chemical processes described by large sparse DAE models, *Fortschritt-Berichte VDI Reihe 3, Verfahrenstechnik*, vol. 613. VDI Verlag, Düsseldorf (1999)
60. Leineweber, D.B., Bauer, I., Bock, H.G., Schlöder, J.P.: An efficient multiple shooting based reduced SQP strategy for large-scale dynamic process optimization. Part I: theoretical aspects. *Computers and Chemical Engineering* 27, 157–166 (2003)
61. Li, W.C., Biegler, L.T.: Multistep, Newton-type control strategies for constrained nonlinear processes. *Chem. Eng. Res. Des.* 67, 562–577 (1989)
62. Liniger, A., Domahidi, A., Morari, M.: Optimization-based autonomous racing of 1:43 scale RC cars. *Optimal Control Applications and Methods* 36(5), 628–647 (2015)
63. Listov, P., Jones, C.: PolyMPC: An efficient and extensible tool for real-time nonlinear model predictive tracking and path following for fast mechatronic systems. *Optimal Control Applications and Methods* 41(2), 709–727 (2020)
64. Löw, S., Obradovic, D., Bottasso, C.: Model predictive control of wind turbine fatigue via online rainflow-counting on stress history and prediction (2020)
65. Maciejowski, J.M.: *Predictive Control with Constraints*. Prentice Hall (2002)
66. MathWorks, T.: The model predictive control toolbox. <https://mathworks.com/products/mpc.html> (2005)
67. Mattingley, J., Boyd, S.: CVXGEN: A code generator for embedded convex optimization. *Optimization and Engineering* pp. 1–27 (2012)
68. Melis, W., Patrinos, P.: C code generation for NMPC. <https://github.com/kulforbes/nmpc-codegen> (2018)
69. Nocedal, J., Wright, S.J.: *Numerical Optimization*, 2 edn. Springer Series in Operations Research and Financial Engineering. Springer (2006)
70. Nurkanović, A., Zanelli, A., Frison, G., Albrecht, S., Diehl, M.: Contraction properties of the advanced step real-time iteration. In: *Proceedings of the IFAC World Congress*, vol. 51 (2020). (accepted)
71. Ohtsuka, T.: A continuation/GMRES method for fast computation of nonlinear receding horizon control. *Automatica* 40(4), 563–574 (2004)
72. Pandala, A.G., Ding, Y., Park, H.W.: qpSWIFT: A real-time sparse quadratic program solver for robotic applications. *IEEE Robotics and Automation Letters* 4(4), 3355–3362 (2019)
73. Patrinos, P., Bemporad, A.: An accelerated dual gradient-projection algorithm for embedded linear model predictive control. *Automatic Control, IEEE Transactions on* 59(1), 18–33 (2014)
74. Qin, S., Badgwell, T.: An overview of industrial model predictive control technology. In: J. Kantor, C. Garcia, B. Carnahan (eds.) *Fifth International Conference on Chemical Process Control – CPC V*, pp. 232–256. American Institute of Chemical Engineers (1996)
75. Quirynen, R., Gros, S., Diehl, M.: Efficient NMPC for nonlinear models with linear subsystems. In: *Proceedings of the IEEE Conference on Decision and Control (CDC)*, pp. 5101–5106 (2013)
76. Quirynen, R., Gros, S., Diehl, M.: Inexact Newton-type optimization with iterated sensitivities. *SIAM Journal on Optimization* 28(1), 74–95 (2018)

77. Quirynen, R., Gros, S., Houska, B., Diehl, M.: Lifted collocation integrators for direct optimal control in ACADO toolkit. *Mathematical Programming Computation* 9(4), 527–571 (2017)
78. Quirynen, R., Houska, B., Diehl, M.: Efficient symmetric Hessian propagation for direct optimal control. *Journal of Process Control* 50, 19–28 (2017)
79. Quirynen, R., Knyazev, A., Di Cairano, S.: Block structured preconditioning within an active-set method for real-time optimal control. In: *Proceedings of the European Control Conference (ECC)* (2018)
80. Rao, C.V., Wright, S.J., Rawlings, J.B.: Application of interior-point methods to model predictive control. *Journal of Optimization Theory and Applications* 99, 723–757 (1998)
81. Rawlings, J.B., Mayne, D.Q., Diehl, M.M.: *Model Predictive Control: Theory, Computation, and Design*, 2nd edn. Nob Hill (2017)
82. Sathya, A., Sopasakis, P., Themelis, A., Parys, R.V., Pipeleers, G., Patrinos, P.: Embedded nonlinear model predictive control for obstacle avoidance using PANOC. In: *Proceedings of the European Control Conference (ECC)* (2018)
83. Schulman, J., Duan, Y., Ho, J., Lee, A., Awwal, I., Bradlow, H., Pan, J., Patil, S., Goldberg, K., Abbeel, P.: Motion planning with sequential convex optimization and convex collision checking. *The International Journal of Robotics Research* 33(9), 1251–1270 (2014)
84. Shukla, H., Khusainov, B., Kerrigan, E., Jones, C.: Software and hardware code generation for predictive control using splitting methods. In: *Proceedings of the IFAC World Congress* (2017)
85. Sopasakis, P., Fresk, E., Patrinos, P.: OpEn: Code generation for embedded nonconvex optimization. In: *IFAC World Congress*
86. Steinbach, M.: *Fast recursive SQP methods for large-scale optimal control problems*. PhD thesis, University of Heidelberg, IWR (1995)
87. Stellato, B., Banjac, G., Goulart, P., Bemporad, A., Boyd, S.: OSQP: An operator splitting solver for quadratic programs. *Mathematical Programming Computation* 12(4), 637–672 (2020)
88. Torrisi, G., Grammatico, S., Smith, R.S., Morari, M.: A projected gradient and constraint linearization method for nonlinear model predictive control. *SIAM Journal on Control and Optimization* 56(3), 1968–1999 (2018)
89. Ullmann, F.: *FiOrdOs: A Matlab toolbox for C-code generation for first order methods*. Master’s thesis, ETH Zurich (2011)
90. Verschueren, R., van Duijkeren, N., Quirynen, R., Diehl, M.: Exploiting convexity in direct optimal control: a sequential convex quadratic programming method. In: *Proceedings of the IEEE Conference on Decision and Control (CDC)* (2016)
91. Verschueren, R., Frison, G., Kouzoupis, D., van Duijkeren, N., Zanelli, A., Quirynen, R., Diehl, M.: Towards a modular software package for embedded optimization. In: *Proceedings of the IFAC Conference on Nonlinear Model Predictive Control (NMPC)* (2018)
92. Verschueren, R., Zanon, M., Quirynen, R., Diehl, M.: A sparsity preserving convexification procedure for indefinite quadratic programs arising in direct optimal control. *SIAM Journal of Optimization* 27(3), 2085–2109 (2017)
93. Wächter, A., Biegler, L.T.: Line Search Filter Methods for Nonlinear Programming: Motivation and Global Convergence. *SIAM Journal on Optimization* 16, 1–31 (2006)
94. Wächter, A., Biegler, L.T.: On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming* 106(1), 25–57 (2006)
95. Wirsching, L., Bock, H.G., Diehl, M.: Fast NMPC of a chain of masses connected by springs. In: *Proceedings of the IEEE International Conference on Control Applications*, Munich, pp. 591–596 (2006)
96. Zanelli, A., Domahidi, A., Jerez, J.L., Morari, M.: FORCES NLP: An efficient implementation of interior-point methods for multistage nonlinear nonconvex programs. *International Journal of Control* (2017)
97. Zanelli, A., Kullick, J., Eldeeb, H., Frison, G., Hackl, C., Diehl, M.: Continuous control set nonlinear model predictive control of reluctance synchronous machines. *IEEE Transactions on Control Systems Technology* (accepted) - <https://arxiv.org/abs/1910.10681> (2019)




Fig. 6 Forces in the springs between the masses in the example of Case Study 1. Replicated from [95].

- 98. Zanelli, A., Quirynen, R., Diehl, M.: Efficient zero-order NMPC with feasibility and stability guarantees. In: Proceedings of the European Control Conference (ECC). Naples, Italy (2019)
- 99. Zometa, P., Kögel, M., Findeisen, R.: muAO-MPC: A free code generation tool for embedded real-time linear model predictive control. In: 2013 American Control Conference, pp. 5320–5325 (2013)

Appendix A - Case study details

Case Study 1: Chain of Masses

System description

The control objective in this example is to stabilize the motion of a chain of $M = 5$ balls with mass m connected by springs to an equilibrium position. The mass on one end of the chain is fixed at $(0, 0, 0)$. The mass on the other end can be freely moved.

Let p_i be the position of mass i , for $i = 1, \dots, M$. The model equations can then be derived as follows. From Hooke’s law, we know that (see Figure 6)

$$F_{i,i+1} = D \left(1 - \frac{L}{\|p_{i+1} - p_i\|} \right) (p_{i+1} - p_i),$$

with each spring having spring constant D and rest length L .

This allows us to write the equations of motion for the middle balls, which read as

$$\ddot{p}_i = \frac{1}{m} (F_{i,i+1} - F_{i-1,i}) + g_z, \quad i = 2, \dots, M - 1,$$

with g_z the gravitational acceleration vector. The control input $u \in \mathbb{R}^3$ directly controls the velocity of the free mass:

$$\dot{p}_M = u.$$

We now introduce a state space formulation with states

$$x^\top = [p_2^\top, p_3^\top, \dots, p_{M-1}^\top, p_M^\top, v_2^\top, v_3^\top, \dots, v_{M-1}^\top] \in \mathbb{R}^{n_x}$$

with $n_x = 3 \cdot (2 \cdot (M - 2) + 1)$, which results in the following ODE:

$$\dot{x} = f(x, u) = \begin{bmatrix} v_2 \\ \vdots \\ v_{M-1} \\ u \\ \frac{1}{m}(F_{2,3} - F_{1,2}) + g_z \\ \vdots \\ \frac{1}{m}(F_{M-1,M} - F_{M-2,M-1}) + g_z \end{bmatrix}. \quad (9)$$

We remark that the only nonlinearity is introduced in the calculation of the forces. The steady state (x_{ss}, u_{ss}) of the system can be found by setting $f(x_{ss}, u_{ss}) = 0$ for any given $p_{M,ss}$. We choose $p_{M,ss} = [7.5, 0, 0]^\top$ for the experiments.

Optimal control problem formulation

In order to stabilize the motion of the chain of masses to the steady state, we propose the following optimal control problem, obtained by performing a multiple shooting discretization of ODE (9):

$$\begin{aligned} \underset{\substack{x_0, \dots, x_N \\ u_0, \dots, u_{N-1}}}{\text{minimize}} \quad & \sum_{k=0}^{N-1} \begin{bmatrix} x_k - x_{\text{ref}} \\ u_k - u_{\text{ref}} \end{bmatrix}^\top \begin{bmatrix} Q & 0 \\ 0 & R \end{bmatrix} \begin{bmatrix} x_k - x_{\text{ref}} \\ u_k - u_{\text{ref}} \end{bmatrix} \\ & + (x_N - x_{\text{ref}})^\top P (x_N - x_{\text{ref}}) \end{aligned} \quad (10)$$

subject to $x_0 = \bar{x}_0$,

$$\begin{aligned} x_{k+1} &= \phi^x(x_k, u_k), \quad k = 0, \dots, N-1, \\ -1 &\leq u_k \leq 1, \quad k = 0, \dots, N-1, \end{aligned}$$

where the initial state \bar{x}_0 is the current estimate of the state vector, $\phi : \mathbb{R}^{n_x} \times \mathbb{R}^3 \rightarrow \mathbb{R}^{n_x}$ is obtained by performing a single RK4 step of length 0.2s on ODE (9). Furthermore, we choose a horizon length $N = 40$, the weighting matrices

$$\begin{aligned} Q &= \text{diag}(\underbrace{0, \dots, 0}_{p_2, \dots, p_{M-1}}, \underbrace{2.5, 2.5, 2.5}_{p_M}, \underbrace{25, \dots, 25}_{v_2, \dots, v_{M-1}}), \\ P &= \text{diag}(\underbrace{0, \dots, 0}_{p_2, \dots, p_{M-1}}, \underbrace{10, 10, 10}_{p_M}, \underbrace{0, \dots, 0}_{v_2, \dots, v_{M-1}}), \\ R &= \text{diag}(0.1, 0.1, 0.1), \end{aligned}$$

the reference control input $u_{\text{ref}} = [0, 0, 0]^\top$ and the state reference

$$x_{\text{ref}} = \left[\underbrace{0, \dots, 0}_{p_2, \dots, p_{M-1}}, \underbrace{7.5, 0, 0}_{p_M}, \underbrace{0, \dots, 0}_{v_2, \dots, v_{M-1}} \right]^\top.$$

Table 7 Design parameters for the chain of masses case study

Quantity	Description	Value
m	mass of one ball	0.1125 kg
D	spring constant	0.4 N/m
L	rest length of the springs	0.1375 m
g_z	gravitational acceleration vector	$[0, 0, -9.81]^\top$ m/s ²
N	horizon length	40
Δt	discretization step	0.2 s
$p_{M,\text{ref}}$	reference position of free ball	$[7.5, 0, 0]^\top$ m




Fig. 7 Schematic illustrating the inverted pendulum on top of a cart.

The design parameters are chosen as in [26] and are summarized in Table 7. Note that we did not introduce path constraints or state bounds, since these are not supported by all solvers that we compare to below.

Case Study 2: Hessian regularization

We control a mass on a rod (a pendulum), balanced on a horizontally moving cart, see Figure 7. The goal is to swing up the pendulum from a stable equilibrium position, namely hanging down.

The dynamics of the cart-pendulum are described by the following ordinary differential equation, where p, v are the horizontal displacement and velocity, respectively, θ is the angle with the vertical and ω the corresponding angular velocity:

$$\dot{p} = v, \tag{11a}$$

$$\dot{\theta} = \omega, \tag{11b}$$

$$\dot{v} = \frac{-ml \sin(\theta) \dot{\theta}^2 + mg \cos(\theta) \sin(\theta) + F}{M + m - m(\cos(\theta))^2}, \tag{11c}$$

$$\dot{\omega} = \frac{-ml \cos(\theta) \sin(\theta) \dot{\theta}^2 + F \cos(\theta) + (M + m)g \sin(\theta)}{l(M + m - m(\cos(\theta))^2)}. \tag{11d}$$

We collect the states in the state vector $x := [p, \theta, v, \omega]^\top$, the control u is the horizontal force F . Transcribing the continuous-time OCP with multiple

shooting gives rise to the following OCP:

$$\underset{\substack{x_0, \dots, x_N, \\ u_0, \dots, u_{N-1}}}{\text{minimize}} \quad \sum_{k=0}^{N-1} \begin{bmatrix} x_k \\ u_k \end{bmatrix}^\top \begin{bmatrix} Q & 0 \\ 0 & R \end{bmatrix} \begin{bmatrix} x_k \\ u_k \end{bmatrix} + x_N^\top Q x_N \quad (12a)$$

$$\text{subject to} \quad x_0 = \bar{x}_0, \quad (12b)$$

$$x_{k+1} = \phi_k^x(x_k, u_k), \quad k = 0, \dots, N-1, \quad (12c)$$

$$-80 \leq u_k \leq 80, \quad k = 0, \dots, N-1, \quad (12d)$$

where ϕ is an RK4 integrator, simulating (11) over one shooting interval. The weight matrices are chosen as

$$Q = \text{diag}([1000, 1000, 0.01, 0.01]), \quad R = 0.01.$$

Because our aim is to swing up the pendulum, we selected strong weights on the position and angle. The other states and the control are assigned a weak penalty in order to avoid too abrupt swing-ups and to favor smooth trajectories. Note that the weighting matrices Q and R are tuning parameters used by the control engineer in the design process in order to obtain a desired behavior. Different choices are therefore equally valid. The initial value is $\bar{x}_0 = [0, \pi, 0, 0]^\top$. We choose $N = 100$ shooting intervals of length 0.01 s.

Case study 3: Hardware-in-the-loop experiments for an engine control application

Two-stage turbocharging gasoline engines are investigated to overcome the drawbacks of conventional (single-stage) turbocharging. The main advantage they offer is a better trade-off between short transient times after load changes and a high specific power. However, the two-stage architecture puts a higher demand on the engine controller, due to the nonlinear nature with cross-couplings in the inputs and the necessity to consider constraints. NMPC has been proposed as a viable control strategy [3].

In Figure 8, a sketch of the two-stage turbocharged engine is depicted. The high-pressure (HP) stage is able to realize fast transients, the low-pressure (LP) stage produces a higher specific power, but with slower dynamics. The control challenge lies in accurately tracking the boost pressure p_{boost} , given the highly nonlinear coupling between both stages.

For reasons of brevity, we directly present the engine model of [3] and refer the interested reader to that work for a derivation. We model the engine with a set of semi-explicit DAEs. The differential states consist of $\Pi_{c,\text{lp}}, \Pi_{c,\text{hp}}$ the pressure ratios between input and output of the compressor in the low pressure and high pressure stage, respectively. The algebraic states are $\Pi_{t,\text{lp}}, \Pi_{t,\text{hp}}$, the pressure ratios on the turbine. The inputs are the wastegate actuation pulse-width modulated signals $u_{\text{wg,lp}}, u_{\text{wg,hp}}$, which take on values between 0% (fully open) and 100% (fully closed).




Fig. 8 Schematic of the two-stage turbocharging concept. Source: [3]

The resulting DAE system reads as

$$\dot{\Pi}_{c,lp} = c_1(\Pi_{t,lp}^{1.5} - \Pi_{t,lp}^{1.25})\sqrt{\Pi_{t,lp}^{-1.5} - \Pi_{t,lp}^{-1.75}} \quad (13)$$

$$- c_2 n_{\text{eng}} \Pi_{c,lp} (\Pi_{c,lp}^{1.29} - \Pi_{c,lp}) \quad (14)$$

$$0 = \Pi_{c,lp} \Pi_{c,hp} \quad (15)$$

$$- \frac{c_3}{n_{\text{eng}}} \sqrt{\Pi_{t,lp}^{0.5} - \Pi_{t,lp}^{0.25}} \left(\sqrt{\Pi_{t,lp}} + c_4 \eta(\Pi_{c,lp} \cdot \Pi_{c,hp}, u_{\text{wg},lp}) \right) \quad (16)$$

$$\dot{\Pi}_{c,hp} = c_5(\Pi_{t,hp}^{1.5} - \Pi_{t,hp}^{1.25})\sqrt{\Pi_{t,hp}^{-1.5} - \Pi_{t,hp}^{-1.75}} \quad (17)$$

$$- c_6 n_{\text{eng}} \Pi_{c,hp} (\Pi_{c,hp}^{1.29} - \Pi_{c,hp}) \quad (18)$$

$$0 = \Pi_{c,lp} \Pi_{c,hp} \quad (19)$$

$$- \frac{c_7}{n_{\text{eng}}} \sqrt{\Pi_{t,hp}^{0.5} - \Pi_{t,hp}^{0.25}} \left(\sqrt{\Pi_{t,hp}} + c_8 (1 - u_{\text{wg},hp}/100) \right), \quad (20)$$

with, additionally, $n_{\text{eng}} = 2000 \text{ min}^{-1}$ the engine speed. We model it as a measured disturbance, in this case a constant. Furthermore, $\eta : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is defined by

$$\eta(u, v) = \gamma_1(u) \cdot \gamma_2(v),$$

with $\gamma_i : \mathbb{R} \rightarrow \mathbb{R}$:

$$\gamma_i(u) = b_{1,i} + b_{2,i} \left(1 + e^{\frac{-u+b_{3,i}}{b_{4,i}}} \right)^{-1}.$$

The values of all model parameters can be found in Table 8.

In order to obtain a smooth control behavior, we include the time derivative of the controls in the optimization formulation, as follows: $\dot{u}_{\text{wg},lp} = d_{u,lp}$, $\dot{u}_{\text{wg},hp} = d_{u,hp}$, and we collect these rates in

$$d = \begin{bmatrix} d_{u,lp} \\ d_{u,hp} \end{bmatrix}.$$

Table 8 PARAMETER VALUES FOR THE TWO-STAGE TURBOCHARGED ENGINE MODEL

Parameter	Unit	Value	Parameter	Unit	Value
c_1	—	25.3	$b_{1,1}$	—	0
c_2	min	0.0034	$b_{2,1}$	—	1
c_3	min^{-1}	7700	$b_{3,1}$	—	1.49
c_4	—	0.6	$b_{4,1}$	—	0.0377
c_5	—	43,6	$b_{1,2}$	—	67.5
c_6	min	0.0092	$b_{2,2}$	—	4.712
c_7	min^{-1}	3600	$b_{3,2}$	—	1
c_8	—	0.9	$b_{4,2}$	—	-1

We then define the vector of differential states, algebraic states and controls, respectively, as follows:

$$x = \begin{bmatrix} \Pi_{c,\text{lp}} \\ \Pi_{c,\text{hp}} \\ u_{\text{wg},\text{lp}} \\ u_{\text{wg},\text{hp}} \end{bmatrix}, \quad z = \begin{bmatrix} \Pi_{t,\text{lp}} \\ \Pi_{t,\text{hp}} \end{bmatrix}, \quad u = \begin{bmatrix} d_{u,\text{lp}} \\ d_{u,\text{hp}} \end{bmatrix}.$$

The values of parameters in OCP (8) are

$$\begin{aligned} y_{r,k} &= [y_{pr,k}; 1.14; 1.54; 50; 50; 0; 0], \\ W &= \text{diag}([10^3, 10^{-3}, 10^{-3}, 10^{-3}, 10^{-3}, 10^{-4}, 10^{-4}]), \\ y_{r,N} &= [y_{pr,N}; 1.14; 1.54; 50; 50], \\ W_N &= \text{diag}([10^3, 10^{-3}, 10^{-3}, 10^{-3}, 10^{-3}]) \end{aligned}$$

with $y_{pr,k}$, $k = 0, \dots, N$ varied as in Figure 5.