

PRESAS: Block-Structured Preconditioning of Iterative Solvers within a Primal Active-Set Method for fast MPC

Quirynen, Rien; Di Cairano, Stefano

TR2020-134 October 07, 2020

Abstract

Model predictive control (MPC) for linear dynamical systems requires solving an optimal control structured quadratic program (QP) at each sampling instant. This paper proposes a primal active-set strategy, called PRESAS, for the efficient solution of such block-sparse QPs, based on a preconditioned iterative solver to compute the search direction in each iteration. Rank-one factorization updates of the preconditioner result in a per-iteration computational complexity of $\mathcal{O}(m^2N)$, where m denotes the number of state and control variables and N the number of control intervals. Three different block-structured preconditioning techniques are presented and their numerical properties are studied further. In addition, an augmented Lagrangian based implementation is proposed to avoid a costly initialization procedure to find a primal feasible starting point. Based on a standalone C code implementation, we illustrate the computational performance of PRESAS against current state of the art QP solvers for multiple linear and nonlinear MPC case studies. We also show that the solver is real-time feasible on a dSPACE MicroAutoBox-II rapid prototyping unit for vehicle control applications, and numerical reliability is illustrated based on experimental results from a testbench of small-scale autonomous vehicles.

Optimal Control Applications & Methods

© 2020 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

ARTICLE TYPE**PRESAS: Block-Structured Preconditioning of Iterative Solvers within a Primal Active-Set Method for fast MPC[†]**

Rien Quirynen* | Stefano Di Cairano

¹Control and Dynamical Systems,
Mitsubishi Electric Research Laboratories,
Massachusetts, USA

Correspondence

*Email: quirynen@merl.com

Present Address201 Broadway, 8th Floor, Cambridge, MA
02139-1955**Abstract**

Model predictive control (MPC) for linear dynamical systems requires solving an optimal control structured quadratic program (QP) at each sampling instant. This paper proposes a primal active-set strategy, called PRESAS, for the efficient solution of such block-sparse QPs, based on a preconditioned iterative solver to compute the search direction in each iteration. Rank-one factorization updates of the preconditioner result in a per-iteration computational complexity of $\mathcal{O}(Nm^2)$, where m denotes the number of state and control variables and N the number of control intervals. Three different block-structured preconditioning techniques are presented and their numerical properties are studied further. In addition, an augmented Lagrangian based implementation is proposed to avoid a costly initialization procedure to find a primal feasible starting point. Based on a standalone C code implementation, we illustrate the computational performance of PRESAS against current state of the art QP solvers for multiple linear and nonlinear MPC case studies. We also show that the solver is real-time feasible on a dSPACE MicroAutoBox-II rapid prototyping unit for vehicle control applications, and numerical reliability is illustrated based on experimental results from a testbench of small-scale autonomous vehicles.

KEYWORDS:

optimization algorithms; model predictive control; preconditioned iterative solvers; active-set methods

1 | INTRODUCTION

Optimization based control and estimation techniques, such as model predictive control (MPC) and moving horizon estimation (MHE), have experienced an increasing amount of interest from industry¹. One of the main advantages of these methods is their systematic way of incorporating system dynamics and inequality constraints. At each sampling instant, MPC solves a multi-stage dynamic optimization problem that minimizes a particular cost function subject to continuity conditions and inequality constraints. A block-sparse quadratic program (QP) structure arises in linear or linear time-varying formulations of predictive control and estimation. A similarly structured quadratic program (QP) forms also the subproblem within a sequential quadratic programming (SQP) method for nonlinear optimal control^{2,3}.

[†]PRESAS: Preconditioned RESidual method within Active-Set strategy for solving constrained linear-quadratic optimal control problems.

In this paper, we are interested in solving the following convex constrained linear-quadratic optimal control problem (OCP)

$$\min_{X,U} \sum_{k=0}^{N-1} \frac{1}{2} \begin{bmatrix} x_k \\ u_k \end{bmatrix}^\top \begin{bmatrix} Q_k & S_k^\top \\ S_k & R_k \end{bmatrix} \begin{bmatrix} x_k \\ u_k \end{bmatrix} + \begin{bmatrix} q_k \\ r_k \end{bmatrix}^\top \begin{bmatrix} x_k \\ u_k \end{bmatrix} + \frac{1}{2} x_N^\top Q_N x_N + q_N^\top x_N \quad (1a)$$

$$\text{s.t.} \quad x_0 = \hat{x}_0, \quad (1b)$$

$$x_{k+1} = a_k + A_k x_k + B_k u_k, \quad k = 0, \dots, N-1, \quad (1c)$$

$$0 \geq d_k + D_k^x x_k + D_k^u u_k, \quad k = 0, \dots, N, \quad (1d)$$

where we define the state vectors as $x_k \in \mathbb{R}^{n_x}$, the control inputs as $u_k \in \mathbb{R}^{n_u}$, the positive semidefinite cost matrices as $Q_k \in \mathbb{R}^{n_x \times n_x}$, $S_k \in \mathbb{R}^{n_u \times n_x}$ and $R_k \in \mathbb{R}^{n_u \times n_u}$, and N is the horizon length. We denote the state and control trajectory, respectively, as $X := [x_0^\top, \dots, x_N^\top]^\top$ and $U := [u_0^\top, \dots, u_{N-1}^\top]^\top$. The constraints include the system dynamics with $A_k \in \mathbb{R}^{n_x \times n_x}$, $B_k \in \mathbb{R}^{n_x \times n_u}$, the inequality constraints with $D_k^x \in \mathbb{R}^{n_{c,k} \times n_x}$, $D_k^u \in \mathbb{R}^{n_{c,k} \times n_u}$ and $\hat{x}_0 \in \mathbb{R}^{n_x}$ denotes the current state estimate.

A real-time implementation of MPC requires the solution of the block-structured QP in (1) within a specified time period, typically on embedded control hardware with limited computational resources and a relatively small amount of available memory⁴. These challenges have sparked considerable efforts in research on embedded optimization algorithms that are tailored to predictive control applications such as described, e.g., in⁵. The following requirements are important to be taken into account when designing or choosing an embedded QP solver for industrial applications of predictive control:

1. scaling of computational complexity and memory requirements with problem dimensions N , n_x and n_u ,
2. deterministic or early termination of solver in real-time applications (obtaining a feasible but suboptimal solution),
3. computational performance on embedded control hardware with relatively limited available resources,
4. warm starting capabilities to reduce solver execution time in receding horizon control applications,
5. numerical reliability and ease of understanding, implementing and maintaining by non-experts,
6. portability of solver code and limited software dependencies for embedded hardware.

Note that the latter list could additionally include infeasibility detection as a requirement for embedded MPC implementations. Tailored infeasibility detection for our proposed PRESAS QP solver is part of future research and remains outside the scope of the present paper. Instead, in what follows, we consider a soft constraint reformulation of all state-dependent inequality constraints, such that the real-time solver can provide a control solution at each time step. Note that input constraints always need to be respected as hard constraints instead. Based on the solution from the embedded optimization algorithm, another component in the overall control system could perform fault detection and recovery.

Most implementations of embedded optimization algorithms that have been successfully applied for real-time optimal control rely on direct linear algebra, e.g., in qpOASES⁶, CVXOPT⁷, FORCES⁸ or HPMPC⁹. However, it is known that iterative methods can result in a better asymptotic complexity when solving the saddle point linear systems arising in second order optimization methods¹⁰. Iterative solvers, such as the minimal residual (MINRES) or the conjugate gradient (CG) method¹¹, are suitable for hardware acceleration of the linear system solution, e.g., using an FPGA for fast embedded applications¹², due to their higher ratio of addition and multiplication operations. The use of an iterative method also allows to trade off between computational time and solution accuracy¹³. For a general linear system, iterative solvers tend to converge rather poorly without preconditioning¹⁰.

Earlier work in^{14,15} studied the use of block-diagonal preconditioners, in combination with numerical techniques that are tailored to an inexact interior-point (IP) framework. More recently, the work in¹⁶ presents an active-set strategy based on a sparse variant of the Newton projection with proportioning (NPP) algorithm and using an augmented Lagrangian-based preconditioner for the MINRES method. Unlike IP methods, an active-set quadratic programming algorithm can considerably benefit from the use of warm or hot-starting techniques to reduce the computational effort when solving a sequence of closely related optimal control problems, see^{17,5,18}. The cost per iteration of an active-set QP solver is computationally cheaper by exploiting low-rank updates of the matrix factorizations when changing the current guess for the active set^{6,19}, compared to standard IP methods¹⁷ or compared to the NPP algorithm in¹⁶. We propose block-structured preconditioning techniques within a primal active-set strategy for real-time optimal control, following our initial investigation in^{20,21}. The resulting Preconditioned RESidual solver within an Active-Set optimization strategy (PRESAS) allows for an initial setup computational complexity of $\mathcal{O}(Nm^3)$ and a per-iteration complexity of $\mathcal{O}(Nm^2)$ for the QP in (1), where m denotes the number of state and control variables in the system.

In the present paper, we provide a complete overview of our recent developments on block-structured preconditioning of iterative solvers with a primal active-set strategy, including block-diagonal preconditioners for MINRES as well as constraint

preconditioning within a projected CG method. We describe theoretical properties for each of the preconditioning techniques, which are validated based on numerical simulation case studies in both single- and double-precision arithmetics. In addition, with respect to our preliminary studies in^{20,21}, in this manuscript we also introduce an improved warm starting of the PRESAS solver, to avoid a costly procedure to find a primal feasible starting point, based on an augmented Lagrangian technique. Using an efficient and self-contained C code implementation, the computational performance of PRESAS is illustrated against state of the art QP solvers. We show that the PRESAS QP solver is real-time feasible on a dSPACE MicroAutoBox-II rapid prototyping unit for vehicle control applications, and numerical reliability is illustrated based on experimental results from a testbench of small-scale autonomous vehicles.

The paper is organized as follows. Section 2 presents the preliminaries on embedded quadratic programming, primal active-set methods and on preconditioning of iterative solvers. Section 3 describes two block-diagonal preconditioners that can be used within a preconditioned minimal residual (PMINRES) solver. Then, Section 4 presents and discusses an optimal control structured constraint preconditioner for the projected preconditioned conjugate gradient (PPCG) method. Two warm-started initialization procedures for the PRESAS solver are presented in Section 5, and the resulting software implementation for embedded MPC applications is described in Section 6. The numerical properties and computational performance of different variants of the proposed PRESAS solver are illustrated in Section 7 based on both linear and nonlinear MPC case studies, including hardware-in-the-loop simulations for vehicle control on a dSPACE MicroAutoBox-II. Section 8 presents experimental results for PRESAS running on a testbench of small-scale autonomous vehicles and Section 9 concludes the paper.

2 | PRELIMINARIES

We assume that the convex QP in (1) has a unique global solution that is non-degenerate. A solution is considered to be degenerate when either the strict complementarity condition or the linear independence constraint qualification (LICQ) does not hold²². In order to have a unique solution to the QP in (1), the Hessian needs to be positive definite on the null-space of the active constraints at the optimal solution. For example, if $S_k = \emptyset$, the weighting matrices are chosen to satisfy $Q_k \geq 0$ and $R_k > 0$ in common MPC formulations, such that the reduced Hessian is automatically positive definite for Eq. (1).

2.1 | Embedded QP Solvers for Optimal Control

There is a general trade-off between solvers that make use of second-order information and require only few but computationally expensive iterations, e.g., qpOASES⁶, versus first-order methods that are of low complexity but may require many more iterations, such as PQP²³, ADMM²⁴ and other gradient or splitting-based methods⁵. In addition, there is an important distinction between optimal control algorithms that target the dense versus the sparse problem formulation. The numerical elimination of the state variables in a condensing routine^{25,26} is typically of a computational complexity $\mathcal{O}(N^2m^3)$ but, unlike nonlinear problems, it can be mostly avoided in linear MPC applications. However, even with an offline preparation of the dense QP formulation, solvers applied to this dense QP will have a runtime complexity of $\mathcal{O}(N^2m^2)$ ¹⁹. Examples of QP solvers to treat the condensed MPC formulation include PQP²³, ADMM²⁴ and qpOASES⁶. Instead, we focus on directly solving the non-condensed MPC formulation with the block sparsity structure in (1), similar to the sparse solvers in FORCES⁸, HPMPC⁹ and qpDUNES²⁷.

It is important to note that many tailored QP algorithms for real-time optimal control rely on strict convexity of the cost function. This enables using a sparsity exploiting dual Newton strategy such as in qpDUNES²⁷, linear algebra routines such as the block-tridiagonal Cholesky factorization of the Schur complement in^{28,29} or the particular Riccati recursion for linear-quadratic control problems in^{30,18}. In the case of a positive semidefinite cost matrix, regularization needs to be applied, followed by an iterative refinement procedure to obtain a solution to the original problem³⁰. This combination of regularization and iterative refinement can be also needed in the presence of ill-conditioned QP matrices. Instead, here, we do not assume strict convexity of the cost function and we propose to use an iterative method to solve each linear KKT system.

2.2 | Primal Feasible Active-Set Method

The basic idea behind active-set methods is to find an optimal active set by iteratively updating a current guess. When fixing the active inequality constraints at the current solution guess, a corresponding structured equality constrained QP is solved to

compute a new search direction

$$\min_{\Delta X, \Delta U} \sum_{k=0}^{N-1} \frac{1}{2} \Delta w_k^\top H_k \Delta w_k + \tilde{h}_k^\top \Delta w_k + \frac{1}{2} \Delta x_N^\top Q_N \Delta x_N + \tilde{q}_N^\top \Delta x_N \quad (2a)$$

$$\text{s.t. } \Delta x_0 = 0, \quad (2b)$$

$$\Delta x_{k+1} = A_k \Delta x_k + B_k \Delta u_k, \quad k = 0, \dots, N-1, \quad (2c)$$

$$0 = D_{k,i}^x \Delta x_k + D_{k,i}^u \Delta u_k, \quad (k, i) \in \mathcal{W}, \quad (2d)$$

where \mathcal{W} denotes the current guess for the active set, i.e., the *working set*. The variables $\Delta w_k := (\Delta x_k, \Delta u_k) = (x_k - \bar{x}_k, u_k - \bar{u}_k)$ are defined for $k = 0, \dots, N-1$ and $\Delta w_N := \Delta x_N = x_N - \bar{x}_N$, such that $\Delta X := [\Delta x_0^\top, \dots, \Delta x_N^\top]^\top$ and $\Delta U := [\Delta u_0^\top, \dots, \Delta u_{N-1}^\top]^\top$, where $\bar{w}_k := (\bar{x}_k, \bar{u}_k)$ denotes the current guess for the optimal solution of the QP in (1). Let us define the Lagrangian for the constrained QP in (1) as follows

$$\begin{aligned} \mathcal{L}(X, U, \lambda_0, \mu_0, \dots, \lambda_N, \mu_N) := & \sum_{k=0}^{N-1} \frac{1}{2} \begin{bmatrix} x_k \\ u_k \end{bmatrix}^\top \begin{bmatrix} Q_k & S_k^\top \\ S_k & R_k \end{bmatrix} \begin{bmatrix} x_k \\ u_k \end{bmatrix} + \begin{bmatrix} q_k \\ r_k \end{bmatrix}^\top \begin{bmatrix} x_k \\ u_k \end{bmatrix} + \frac{1}{2} x_N^\top Q_N x_N + q_N^\top x_N + \lambda_0^\top (\hat{x}_0 - x_0) \\ & + \sum_{k=0}^{N-1} \lambda_{k+1}^\top (a_k + A_k x_k + B_k u_k - x_{k+1}) + \sum_{k=0}^N \mu_k^\top (d_k + D_k^x x_k + D_k^u u_k), \end{aligned} \quad (3)$$

where λ_k and μ_k denote the Lagrange multipliers, respectively, for equality and inequality constraints at each stage $k = 0, \dots, N$.

The Hessian matrices $H_k := \begin{bmatrix} Q_k & S_k^\top \\ S_k & R_k \end{bmatrix}$ and Lagrangian gradients $\tilde{h}_k := H_k \bar{w}_k + \begin{bmatrix} q_k - \lambda_k \\ r_k \end{bmatrix} + \begin{bmatrix} A_k^\top \\ B_k^\top \end{bmatrix} \lambda_{k+1} + \sum_{i \in \mathcal{W}_k} \mu_{k,i} \begin{bmatrix} D_{k,i}^{x^\top} \\ D_{k,i}^{u^\top} \end{bmatrix}$

are defined for $k = 0, \dots, N-1$ and $\tilde{q}_N := Q_N \bar{x}_N + q_N - \lambda_N + \sum_{i \in \mathcal{W}_N} \mu_{N,i} \begin{bmatrix} D_{N,i}^{x^\top} \\ D_{N,i}^{u^\top} \end{bmatrix}$.

The equality constrained QP in (2) results in the search direction $\bar{w}_k + \alpha \Delta w_k$ for which all constraints in the set \mathcal{W} remain satisfied, regardless of the value $\alpha \geq 0$. A distinction should be made between primal, dual and primal-dual active-set methods²². In addition, parametric methods have been proposed⁶ in order to exploit the parametric aspect within a homotopy framework. The active-set identification suffers from a combinatorial complexity and is therefore *NP hard* in the worst case¹⁷. However, active-set methods have been successfully used in many real-time control applications⁵. In this work, we consider the general framework of primal feasible active-set methods as described in²², see Algorithm 1, and which generally provides the following numerical properties and advantages over alternative active-set optimization strategies:

1. does not require strong duality and is relatively easy to implement compared to, e.g., a primal-dual method,
2. can be terminated at any time point and still provides a feasible, yet suboptimal, solution to the OCP,
3. automatically maintains the linear independence of the constraint gradients in its working set²².

The latter property is very important in practice since it avoids the need for computing the maximal linearly independent subset of the current active set as described in⁶ for dense QP problems. In case of sparse optimal control problems, the matrix factorization update techniques for active-set methods become relatively complicated, as discussed in¹⁹. Also, it becomes nontrivial or impossible to efficiently maintain the linear independence of the working set in combination with a direct exploitation of the block sparsity structure in optimal control problems. Therefore, we focus on a primal feasible active-set method where such linear independence is automatically preserved. The work in^{31,32} describes an approach for complexity certification to compute a partitioning of the parameter space and corresponding bounds on the number of iterations in primal active-set methods, which could be applied directly to our proposed optimization algorithm for MPC. Procedures for handling degeneracy and cycling in primal active-set methods can be found in²², but such issues are not common in practical MPC problems. Note that the primal active-set method in Algorithm 1 adds (step 6) or removes (step 11) exactly one constraint in the current working set at each iteration. Multiple constraints with negative Lagrange multiplier values could instead be removed in step 11, but we restrict to the standard implementation in Algorithm 1 for the remainder of this article.

2.3 | Block-Structured Linear System Solution

At each iteration of the active-set method, the solution of the equality constrained QP in (2) corresponds to solving a saddle point linear system that represents the first order necessary conditions of optimality, i.e., the Karush-Kuhn-Tucker (KKT) conditions.

The block-structured saddle point linear system can be written in compact form as follows

$$\begin{bmatrix} \mathcal{H} & \mathcal{A}^\top \\ \mathcal{A} & \mathbb{0} \end{bmatrix} \begin{bmatrix} \Delta y \\ \Delta v \end{bmatrix} = \begin{bmatrix} \bar{b} \\ \mathbb{0} \end{bmatrix} \quad \text{or} \quad \mathcal{K} z = \bar{b}, \quad (4)$$

where the primal variables $\Delta y := [\Delta w_0^\top, \dots, \Delta w_N^\top]^\top$ and the Lagrange multipliers $\Delta v := [\Delta \lambda_0^\top, \Delta \mu_0^\top, \dots, \Delta \lambda_N^\top, \Delta \mu_N^\top]^\top$ are defined. Similarly, the right-hand side vector is defined as $b := -[\tilde{h}_0^\top, \dots, \tilde{h}_{N-1}^\top, \tilde{q}_N^\top]^\top$ based on the Lagrangian gradient vectors. The linear system in (4) has a particular sparsity structure because \mathcal{H} is the block-diagonal Hessian matrix and the constraint matrix \mathcal{A} reads as

$$\mathcal{H} = \begin{bmatrix} H_0 & & & \\ & H_1 & & \\ & & \ddots & \\ & & & Q_N \end{bmatrix}, \quad \mathcal{A} = \begin{bmatrix} -\mathbb{1} & \mathbb{0} & & & \\ E_0^x & E_0^u & & & \\ A_0 & B_0 & -\mathbb{1} & \mathbb{0} & \\ & & & \ddots & \\ & & & & \ddots \end{bmatrix} = \begin{bmatrix} [-\mathbb{1} \ \mathbb{0}] \\ E_0 \\ C_0 \\ [-\mathbb{1} \ \mathbb{0}] \\ \ddots \end{bmatrix}, \quad (5)$$

where $\mathbb{1}$ and $\mathbb{0}$ denote, respectively, the identity and zero matrix, $E_k^x := [D_{k,i}^x]_{i \in \mathcal{W}_k}$ and $E_k^u := [D_{k,i}^u]_{i \in \mathcal{W}_k}$ denote the active inequality constraints for each interval k , corresponding to the working set in \mathcal{W} , and we define the block matrices $C_k := [A_k \ B_k]$ and $E_k := [E_k^x \ E_k^u]$ for the sake of simplifying the notation.

Algorithm 1 Primal Feasible Active-Set Method for Quadratic Programming.

Input: Primal feasible starting point y_0 and working set \mathcal{W}_0 .

- 1: **for** iter = 1, 2, ..., max **do**
- 2: Solve the QP in (2), i.e., the linear system in (4), to compute the search direction $(\Delta y, \Delta v)$.
- 3: $\alpha \leftarrow \min_{(k,j) \notin \mathcal{W}} -\frac{d_{k,j} + D_{k,j} \bar{w}_k}{D_{k,j} \Delta w_k}$. (step length computation)
- 4: **if** $\alpha < 1$ **then**
- 5: $(k, j) \leftarrow \arg \min_{(k,j) \notin \mathcal{W}} -\frac{d_{k,j} + D_{k,j} \bar{w}_k}{D_{k,j} \Delta w_k}$. (most blocking constraint)
- 6: $\mathcal{W} \leftarrow \mathcal{W} \cup \{(k, j)\}$. (add constraint to working set)
- 7: $y^* \leftarrow y + \alpha \Delta y$ and $v^* \leftarrow v + \alpha \Delta v$.
- 8: **else**
- 9: $y^* \leftarrow y + \Delta y$ and $v^* \leftarrow v + \Delta v$.
- 10: **if** $\exists (k, j) \in \mathcal{W} : \mu_{k,j}^* < 0$ **then**
- 11: $\mathcal{W} \leftarrow \mathcal{W} \setminus \{\arg \min_{(k,j) \in \mathcal{W}} \mu_{k,j}^*\}$. (remove from working set)
- 12: **else**
- 13: **stop** (QP solution found)

Output: Optimal solution y^* and v^* for the QP in (1).

2.4 | Preconditioning of Iterative Solvers

In Eq. (4), the matrix \mathcal{A} has full rank and \mathcal{H} is symmetric and positive semidefinite. Unlike prior work on embedded optimization algorithms for optimal control based on direct linear algebra routines, e.g.,^{30,19,29,18}, we propose the use of iterative solvers as discussed for general saddle point linear systems in^{10,33}. However, preconditioning is necessary for the good performance of iterative solvers¹³. Standard preconditioning results in a modified linear system $\mathcal{T}^{-1} \mathcal{K} z = \mathcal{T}^{-1} b$ where \mathcal{T} is the preconditioner. Instead, based on a positive definite preconditioning matrix $\mathcal{T} = LL^\top$, symmetric preconditioning $L^{-1} \mathcal{K} L^{-\top}$ can be performed, in which the preconditioner \mathcal{T} should be designed such that

1. computations with the operator $\mathcal{T} = LL^\top$ are cheaper than solving the original saddle point linear system in (4),
2. the preconditioned matrix $L^{-1} \mathcal{K} L^{-\top}$ approximates the identity or its eigenvalues are tightly clustered¹¹.

An overview on algebraic and application-specific preconditioners can be found in¹⁰. Here, we focus on two block-diagonal preconditioning techniques and one constraint preconditioner that is tailored to the projected conjugate gradient method. We aim to utilize warm- or hot-starting capabilities of active-set methods and exploit the rank-one property of active-set changes in the updating procedure for the factorization of the structured preconditioner from one active-set iteration to the next.

computational complexity for the block-tridiagonal Cholesky factorization of the matrix $\mathcal{A}\tilde{\mathcal{H}}^{-1}\mathcal{A}^\top$, as part of the preconditioner in (8), is $\mathcal{O}(Nm^3)$ in which $m = n_x + n_u$. Similar to before, a rank-one factorization update can be computed, resulting in a computational complexity of $\mathcal{O}(Nm^2)$ at each iteration of the active-set method.

In case that a particular Hessian block matrix H_k is positive semidefinite, a positive definite approximation $\tilde{H}_k = \begin{bmatrix} \tilde{Q}_k & \tilde{S}_k^\top \\ \tilde{S}_k & \tilde{R}_k \end{bmatrix} \succ 0$ needs to be computed. One can apply an *on-the-fly* regularization in the form of a modified Cholesky factorization, e.g., as in^{27,22}, or alternatively the structure-exploiting regularization technique tailored to optimal control that was recently proposed in³⁴. Here, we consider a simple but standard regularization procedure of the form $\tilde{H}_k = H_k + \epsilon \mathbb{1}$, where the value for $\epsilon > 0$ is chosen sufficiently small such that $\tilde{H}_k \approx H_k$ forms a good Hessian approximation, but large enough such that $\tilde{H}_k \succ 0$ is positive definite and leads to good numerical conditioning of \mathcal{T}_s in (9).

3.3 | Preconditioned MINRES Algorithm

Both the Schur-complement based \mathcal{T}_s and the augmented Lagrangian preconditioner \mathcal{T}_a are symmetric positive definite, given any positive value for $\epsilon > 0$ or $\gamma > 0$. Therefore, these preconditioning techniques can be used within the preconditioned minimal residual (PMINRES) algorithm as described in¹¹. PMINRES is based on the three-term recurrence in the Lanczos iteration for symmetric matrices and specifically requires a positive definite preconditioner¹⁰. On the other hand, the generalized minimal residual (GMRES) algorithm is based on the Arnoldi iteration, for which the computational cost generally grows with each iteration, but it does not need a positive definite preconditioner. Here, we focus on using PMINRES in combination with either of the block-diagonal preconditioners in Eq. (6) or (8).

As discussed in detail by³⁵, as the value of $\gamma > 0$ increases for the augmented Lagrangian preconditioner $\mathcal{T}_a = L_a L_a^\top$, the eigenvalues of the preconditioned matrix $L_a^{-1} \mathcal{K} L_a^{-\top}$ become more and more tightly clustered around ± 1 . For a sufficiently large value of γ , MINRES can therefore converge within two iterations in the ideal setting, but choosing γ too large may result in ill-conditioning of the matrix \mathcal{T}_a . For the Schur-complement based preconditioner $\mathcal{T}_s = L_s L_s^\top$, when $\tilde{\mathcal{H}} = \mathcal{H}$ is invertible, the preconditioned matrix has three distinct eigenvalues 1 , $\frac{1}{2}(1 + \sqrt{5})$ and $\frac{1}{2}(1 - \sqrt{5})$, and therefore MINRES converges within three iterations³⁶. The eigenvalues become clustered when, e.g., regularization is needed to make $\tilde{\mathcal{H}} \approx \mathcal{H}$ positive definite^{10,36}. Similar to before, for a smaller value of $\epsilon > 0$ in $\tilde{\mathcal{H}} = \mathcal{H} + \epsilon \mathbb{1}$, the eigenvalues of $L_s^{-1} \mathcal{K} L_s^{-\top}$ become more tightly clustered. Yet, when choosing the value of $\epsilon > 0$ too small, ill-conditioning of the matrix \mathcal{T}_s may cause numerical issues. The effect of these choices on the overall performance of the algorithm will be illustrated in the case studies of Section 7.

4 | PROJECTED PRECONDITIONED CONJUGATE GRADIENT METHOD

The parameters, γ or ϵ , in the above mentioned preconditioners can be difficult to choose, especially when using lower precision arithmetics or in the presence of ill-conditioned QP matrices. Therefore, in what follows, we present a preconditioned iterative solver for which the performance is less dependent on a careful choice for the regularization parameter. We propose to use conjugate gradient (CG) iterations in the null space of the active inequality constraints, based on a particular projection operator that preserves the block-structured sparsity in optimal control.

4.1 | Conjugate Gradient for Reduced Linear System

The equality constrained QP in (2) could be solved by eliminating the constraints and solving the resulting reduced problem formulation. Next, we briefly review the issues that arise when applying the CG method to such a reduced formulation of the linear KKT system. We define the matrix \mathcal{Z} as a basis for the null space of the constraint matrix \mathcal{A} such that $\mathcal{A}\mathcal{Z} = 0$. Given the linear system in (4), the reduced formulation then reads as

$$\mathcal{Z}^\top \mathcal{H} \mathcal{Z} \Delta y_z = \mathcal{Z}^\top b, \quad (10)$$

where $\Delta y = \mathcal{Z} \Delta y_z$, i.e., the solution vector Δy is in the null space of the active constraint matrix \mathcal{A} .

One could directly apply the conjugate gradient (CG) method with standard preconditioning techniques to the reduced linear system in (10) as described in^{37,38}. This is generally not advisable for optimal control problems because the block-structured sparsity is destroyed, similar to the use of a condensing routine²⁵, resulting in expensive CG iterations of complexity $\mathcal{O}(N^2 m^2)$ instead of the desired runtime complexity of $\mathcal{O}(Nm^2)$. In addition, forming a null space basis matrix \mathcal{Z} , computing the reduced

Hessian $\mathcal{Z}^\top \mathcal{H} \mathcal{Z}$ and constructing and applying a preconditioner for this reduced Hessian is computationally expensive³⁸, in general. Sparsity exploiting update procedures for the null space basis matrix and for the factorization in a Schur complement step of an active-set strategy have been proposed in¹⁹, but they are relatively complex to implement.

4.2 | Projected Preconditioned Conjugate Gradient (PPCG)

Instead of reducing to the linear system in Eq. (10), we propose to apply a projection operator \mathcal{P} , onto the null space of \mathcal{A} , to the linear KKT system in Eq. (4):

$$\mathcal{P} \mathcal{H} \Delta y = \mathcal{P} b, \quad (11)$$

where $\mathcal{P} \mathcal{A}^\top = 0$ ³⁹. One possible definition of the projection operator is $\mathcal{P} := \mathcal{Z} (\mathcal{Z}^\top \mathcal{Z})^{-1} \mathcal{Z}^\top$, but it requires again the computation of a basis matrix \mathcal{Z} for the null space of \mathcal{A} . The idea behind the projected preconditioned conjugate gradient (PPCG) method from^{37,40,38} is to use the alternative projection operator

$$\mathcal{P}_{\mathcal{A}:\tilde{\mathcal{H}}} := \mathbb{1} - \mathcal{A}^\top (\mathcal{A} \tilde{\mathcal{H}}^{-1} \mathcal{A}^\top)^{-1} \mathcal{A} \tilde{\mathcal{H}}^{-1}, \quad (12)$$

where $\mathbb{1}$ denotes the identity matrix. The projection in (12) does not require any basis matrix \mathcal{Z} and preserves the sparsity of a block-structured problem, due to the block-tridiagonal sparsity structure of the positive definite matrix $\mathcal{A} \tilde{\mathcal{H}}^{-1} \mathcal{A}^\top$ and the block-diagonal structure of the positive definite matrix $\tilde{\mathcal{H}}$, as discussed in the next subsection. The projection operator in (12) is defined for a particular preconditioner $\tilde{\mathcal{H}} > 0$ and $\tilde{\mathcal{H}} \approx \mathcal{H}$ that can be applied directly to Eq. (11), resulting in

$$\tilde{\mathcal{H}}^{-1} \mathcal{P}_{\mathcal{A}:\tilde{\mathcal{H}}} \mathcal{H} \Delta y = \tilde{\mathcal{H}}^{-1} \mathcal{P}_{\mathcal{A}:\tilde{\mathcal{H}}} b. \quad (13)$$

One can then apply the conjugate gradient method to the projected preconditioned linear system in (13).

The operation $(\tilde{\mathcal{H}}^{-1} \mathcal{P}_{\mathcal{A}:\tilde{\mathcal{H}}}) b_1$ in Eq. (13) can alternatively be computed by solving the linear system

$$\begin{bmatrix} \tilde{\mathcal{H}} & \mathcal{A}^\top \\ \mathcal{A} & 0 \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} b_1 \\ 0 \end{bmatrix}, \quad \text{or} \quad \mathcal{T}_c z = \bar{b}_1, \quad (14)$$

for a certain vector b_1 . The matrix \mathcal{T}_c is also referred to as a *constraint preconditioner* for the KKT system in (4). The solution of the linear system (14) corresponds to

$$z_1 := \tilde{\mathcal{H}}^{-1} (b_1 - \mathcal{A}^\top z_2), \quad z_2 := (\mathcal{A} \tilde{\mathcal{H}}^{-1} \mathcal{A}^\top)^{-1} \mathcal{A} \tilde{\mathcal{H}}^{-1} b_1, \quad (15)$$

such that $z_1 = (\tilde{\mathcal{H}}^{-1} \mathcal{P}_{\mathcal{A}:\tilde{\mathcal{H}}}) b_1$ holds. Furthermore, we use the improved variant of the PPCG method based on the residual update strategy in^{40,38}, as described in Algorithm 2, in order to avoid significant round-off errors in the range of \mathcal{A}^\top . An additional iterative refinement step could also be used to improve the numerical stability of the solver.

4.3 | Block-structured Sparse Constraint Preconditioner

The standard block-tridiagonal sparsity structure of the positive definite matrix $\mathcal{A} \tilde{\mathcal{H}}^{-1} \mathcal{A}^\top$ as well as the block-diagonal structure of the positive definite matrix $\tilde{\mathcal{H}}$ can be exploited in Eq. (15) for the implementation of PRESAS in combination with the PPCG method for fast MPC applications. In addition, the approximate Hessian matrix $\tilde{\mathcal{H}}$ in the preconditioner is often chosen to be diagonal, in practice. As described earlier in Section 3, the block-tridiagonal Cholesky factorization for the matrix $\mathcal{A} \tilde{\mathcal{H}}^{-1} \mathcal{A}^\top$ can be efficiently maintained based on a rank-one factorization update for each active-set change. Similar to the techniques in Section 3, this variant of the solver requires an initial setup computational complexity of $\mathcal{O}(Nm^3)$ and a per-iteration complexity of $\mathcal{O}(Nm^2)$ for solving the optimal control structured QP in (1).

The quality of the constraint preconditioner is defined by the accuracy of the Hessian approximation $\tilde{\mathcal{H}} \approx \mathcal{H}$ such that $\tilde{\mathcal{H}} > 0$. However, unlike the case for the Schur complement type block-diagonal preconditioner, the eigenvalues of the preconditioned matrix can be shown to be

$$\{1\} \cup \sigma \left((\mathcal{Z}^\top \tilde{\mathcal{H}} \mathcal{Z})^{-1} \mathcal{Z}^\top \mathcal{H} \mathcal{Z} \right), \quad (16)$$

where $\sigma(\cdot)$ denotes the spectrum of a matrix. As described in³⁸, typical choices are $\tilde{\mathcal{H}} = \mathbb{1}$ or $\tilde{\mathcal{H}} = \text{diag}(\mathcal{H})$. Here, we use again a regularized Hessian approximation of the form $\tilde{\mathcal{H}} = \mathcal{H} + \epsilon \mathbb{1}$, where $\epsilon > 0$ is chosen sufficiently small without causing ill-conditioning of the constraint preconditioner.

Two conclusions can be made straightforwardly from the expression for the eigenvalues of the preconditioned matrix in Eq. (16). First, all eigenvalues are equal to 1 in case that $\mathcal{H} = \tilde{\mathcal{H}} > 0$, i.e., only one CG iteration is needed. Second, all eigenvalues

Algorithm 2 Projected Preconditioned Conjugate Gradient (PPCG) method with the residual update strategy.

Input: Matrices \mathcal{H} , \mathcal{A} , Hessian approximation $\tilde{\mathcal{H}}$ and vector b in (4).

```

1:  $\Delta y \leftarrow 0$  and  $r \leftarrow \mathcal{H} \Delta y - b$ .
2: Solve  $\begin{bmatrix} \tilde{\mathcal{H}} & \mathcal{A}^\top \\ \mathcal{A} & 0 \end{bmatrix} \begin{bmatrix} g \\ v \end{bmatrix} = \begin{bmatrix} r \\ 0 \end{bmatrix}$  such as in Eq. (15).
3:  $p \leftarrow -g$ ,  $\Delta v \leftarrow -v$  and  $r \leftarrow r - \mathcal{A}^\top v$ .
4: for  $n = 1, \dots, n_{\max}$  do
5:    $\alpha \leftarrow \frac{r^\top g}{p^\top \tilde{\mathcal{H}} p}$ .
6:    $\Delta y \leftarrow \Delta y + \alpha p$ .
7:    $r^+ \leftarrow r + \alpha \mathcal{H} p$ .
8:   if  $\|r^+\| < \epsilon_{\text{tol}}$  then
9:      $\Delta y^* \leftarrow \Delta y$  and  $\Delta v^* \leftarrow \Delta v$ .
10:    stop (solution found)
11:   else
12:     Solve  $\begin{bmatrix} \tilde{\mathcal{H}} & \mathcal{A}^\top \\ \mathcal{A} & 0 \end{bmatrix} \begin{bmatrix} g^+ \\ v^+ \end{bmatrix} = \begin{bmatrix} r^+ \\ 0 \end{bmatrix}$  such as in Eq. (15).
13:      $\Delta v \leftarrow \Delta v - v^+$ .
14:      $\beta \leftarrow \frac{r^{+\top} g^+}{r^{+\top} g}$ .
15:      $p \leftarrow -g^+ + \beta p$ .
16:      $g \leftarrow g^+$ .
17:      $r \leftarrow r^+ - \mathcal{A}^\top v^+$ .

```

Output: Solution vectors Δy^* and Δv^* for linear system in Eq. (4).

remain equal to 1 also when an augmented Lagrangian type regularization is applied to the Hessian, i.e., $\tilde{\mathcal{H}} = \mathcal{H} + \mathcal{A}^\top \Gamma \mathcal{A} > 0$ such that the reduced Hessian reads $\mathcal{Z}^\top \tilde{\mathcal{H}} \mathcal{Z} = \mathcal{Z}^\top \mathcal{H} \mathcal{Z}$ as in⁴⁰. The latter is very common for MPC applications, e.g., when performing a slack reformulation of inequality constraints based on an exact L1 penalty. The full Hessian matrix will be positive semidefinite as a result of such slack reformulation, but a regularization of the slack variables in the constraint preconditioner does not change the reduced Hessian and therefore it does not affect the convergence of the PPCG method. In practice, the number of iterations can be larger than one for both of these cases because of numerical round-off errors that are caused by ill-conditioning and/or the usage of low-precision arithmetics.

5 | WARM-STARTED INITIALIZATION OF PRESAS FOR PREDICTIVE CONTROL

As indicated in Algorithm 1, a primal active-set method requires an initial point that is primal feasible. In the general case of a constrained quadratic program, this is a nontrivial task that corresponds to a *Phase I* procedure as described, for example, in^{41,22}. Here, we describe two approaches that are tailored to the solution of constrained optimal control problems, in order to allow an infeasible start of the primal active-set method and therefore avoid the need for an explicit *Phase I* procedure.

5.1 | Forward Simulation of Dynamics and Slack Reformulation

When solving the parametric OCP in (1) within receding horizon based control or estimation, it becomes relatively easy to satisfy the initial value condition and the continuity constraints based on a forward simulation of the state variables using a shifted version of the control trajectory. In addition, one may introduce slack variables in order to always satisfy all state-dependent inequality constraints in (1d). The first approach is therefore based on a penalty method, similar in spirit to that described in²².

Let us reformulate the structured QP obtained from optimal control problem (1) as follows:

$$\min_{x, U, S} \sum_{k=0}^{N-1} \frac{1}{2} \begin{bmatrix} x_k \\ u_k \\ s_k \end{bmatrix}^\top \begin{bmatrix} Q_k & S_k^\top & \\ S_k & R_k & \\ & & M_k \end{bmatrix} \begin{bmatrix} x_k \\ u_k \\ s_k \end{bmatrix} + \begin{bmatrix} q_k \\ r_k \\ m_k \end{bmatrix}^\top \begin{bmatrix} x_k \\ u_k \\ s_k \end{bmatrix} + \frac{1}{2} \begin{bmatrix} x_N \\ s_N \end{bmatrix}^\top \begin{bmatrix} Q_N & \\ & M_N \end{bmatrix} \begin{bmatrix} x_N \\ s_N \end{bmatrix} + \begin{bmatrix} q_N \\ m_N \end{bmatrix}^\top \begin{bmatrix} x_N \\ s_N \end{bmatrix} \quad (17a)$$

$$\text{s.t. } x_0 = \hat{x}_0, \quad (17b)$$

$$x_{k+1} = a_k + A_k x_k + B_k u_k, \quad k = 0, \dots, N-1, \quad (17c)$$

$$s_k \geq d_k + D_k^x x_k + D_k^u u_k, \quad k = 0, \dots, N, \quad (17d)$$

$$s_k \geq 0, \quad k = 0, \dots, N, \quad (17e)$$

where $M_k \geq 0$ and $m_k \geq 0$, respectively, represent the L2 and exact L1 penalty for each of the slack variables $S := [s_0, \dots, s_N]^\top$. It can be shown that $s_k = 0$ for each $k = 0, \dots, N$ in the solution of the soft-constrained QP (17), whenever a feasible solution exists for the QP in (1) and for a sufficiently large value of $m_k > 0$ in the exact L1 penalty⁴¹.

The resulting solver is easy to implement, because there is no explicit *Phase I* procedure needed for initialization and a primal active-set method automatically maintains the linear independence of the constraint gradients in its working set²². However, the forward simulation procedure to satisfy the equality constraints in eqs. (17b)-(17c) may lead to a poor initial guess for the primal active-set method in case of unstable system dynamics and/or large changes of the initial state value in (17b), after shifting the state and control trajectory from one control time step to the next.

5.2 | Augmented Lagrangian Method for Initial Value Condition

In order to avoid the potential numerical instability in the forward simulation procedure due to unstable system dynamics and/or large changes of the initial state value, we propose to enforce the initial state value condition of Eq. (17b) in the form of an augmented Lagrangian method as described in²². For the soft-constrained optimal control structured problem in (17), we remove the initial state value constraint and instead compute the solution to the following QP

$$\min_{x, U, S} \sum_{k=0}^{N-1} \frac{1}{2} \begin{bmatrix} x_k \\ u_k \\ s_k \end{bmatrix}^\top \begin{bmatrix} Q_k & S_k^\top & \\ S_k & R_k & \\ & & M_k \end{bmatrix} \begin{bmatrix} x_k \\ u_k \\ s_k \end{bmatrix} + \begin{bmatrix} q_k \\ r_k \\ m_k \end{bmatrix}^\top \begin{bmatrix} x_k \\ u_k \\ s_k \end{bmatrix} + \frac{1}{2} \begin{bmatrix} x_N \\ s_N \end{bmatrix}^\top \begin{bmatrix} Q_N & \\ & M_N \end{bmatrix} \begin{bmatrix} x_N \\ s_N \end{bmatrix} + \begin{bmatrix} q_N \\ m_N \end{bmatrix}^\top \begin{bmatrix} x_N \\ s_N \end{bmatrix} \quad (18a)$$

$$+ \frac{1}{2} \|x_0 - \hat{x}_0\|_{P^i}^2 + \lambda_0^{i^\top} (\hat{x}_0 - x_0) \quad (18b)$$

$$\text{s.t. } x_{k+1} = a_k + A_k x_k + B_k u_k, \quad k = 0, \dots, N-1, \quad (18c)$$

$$s_k \geq d_k + D_k^x x_k + D_k^u u_k, \quad k = 0, \dots, N, \quad (18d)$$

$$s_k \geq 0, \quad k = 0, \dots, N, \quad (18e)$$

at each iteration $i = 0, \dots$, where λ_0^i and $P^i := \text{diag}(\rho^i)$, respectively, denote the Lagrange multipliers and the quadratic penalty for the initial value condition of Eq. (17b). As discussed in²², after solving the QP (18), the Lagrange multiplier values for the initial state value condition can be updated by

$$\lambda_0^{i+1} = \lambda_0^i + P^i (\hat{x}_0 - x_0). \quad (19)$$

An increasingly large value can be chosen for the quadratic penalty $P^{i+1} \geq P^i$, even though convergence of the augmented Lagrangian method can be assured even without increasing the value for $P := \text{diag}(\rho)$ ²². In addition, each quadratic penalty can be updated individually for each of the initial state vector components.

Algorithm 3 provides a detailed description of the augmented Lagrangian type implementation for effective warm starting of the proposed primal active-set strategy in the PRESAS QP solver, which will be referred to as PRESAS-WSA. Note that the QP solution in each iteration is used as the warm-started initialization of PRESAS in the next augmented Lagrangian type iteration. In addition, the block-structured Cholesky factorization for the preconditioning matrix can be reused from one iteration to the next, resulting in a *hot start* for each call to the PRESAS QP solver within the augmented Lagrangian method of Algorithm 3. Rank-one factorization updates are needed for each of the individual changes in the quadratic penalty values, corresponding to a rank-one update for the Hessian matrix of the QP subproblem (18). However, as mentioned also earlier, when using a reverse block-tridiagonal Cholesky factorization procedure²⁷, the computational cost per rank-one update in the first stage of the control horizon becomes $\mathcal{O}(m^2)$ instead of the computational complexity of $\mathcal{O}(Nm^2)$ for the complete factorization update.

Algorithm 3 Warm-starting of PRESAS: Augmented Lagrangian Method for Initial State Value (PRESAS-WSA).

Input: Initial solution guess X^0, U^0, S^0 and initial values $P^0 = \text{diag}(\rho^0)$ and λ^0 .

- 1: Forward simulation of system dynamics to ensure satisfaction of Eq. (17c).
- 2: Increase slack variables $s_k \geq 0$ for $k = 0, \dots, N$ to satisfy inequality constraints (17d).
- 3: **while** $\|x_0 - \hat{x}_0\| > \epsilon_{\text{tol}}$ **do**
- 4: Solve block-structured QP subproblem in (18) using Algorithm 4 (PRESAS), with reuse of Cholesky factorization.
- 5: Update Lagrange multiplier values: $\lambda_0^{i+1} \leftarrow \lambda_0^i + P^i (x_0 - \hat{x}_0)$.
- 6: **for** $j = 1, \dots, n_x$ **do**
- 7: **if** $|x_{0,j} - \hat{x}_{0,j}| > \epsilon_{\text{tol}}$ AND $\rho_j^i < \rho_{\text{max}}$ **then**
- 8: Update quadratic penalty value: $\rho_j^{i+1} \geq \rho_j^i$ for initial state variable $x_{0,j}$.
- 9: **if** $\rho_j^{i+1} > \rho_j^i$ **then**
- 10: Rank-1 Cholesky factorization update for preconditioner in PRESAS.
- 11: **end while**

Output: Solution variables X^*, U^* and S^* for the soft-constrained QP in (17).

The computational performance of PRESAS, using either the penalty or the augmented Lagrangian method, is illustrated based on closed-loop numerical simulations for case studies of linear and nonlinear MPC in Section 7.

6 | SOFTWARE IMPLEMENTATION OF PRESAS FOR EMBEDDED MPC APPLICATIONS

The proposed PRESAS solver is described in Algorithm 4, based on the primal active-set strategy in Algorithm 1 and the PFCG method in Algorithm 2, in combination with structure exploiting factorization updates for the block-sparse preconditioner. The solver adheres to all embedded optimal control requirements (1-6) that were stated in the introduction. Namely, PRESAS enjoys the preferred computational complexity of $\mathcal{O}(Nm^2)$ per iteration based on its block sparsity structure exploitation and the use of iterative linear algebra routines. For being an active-set method with tailored optimal control type structure exploitation, the proposed approach is considerably easier to implement than prior work such as¹⁹.

6.1 | Self-contained C code Implementation for Embedded Applications

In what follows, we illustrate the numerical properties and the computational performance of the proposed primal active-set solver for fast MPC applications, based on a preliminary C code implementation. In order to simplify its use for industrial applications of optimization based control and estimation, the software implementation does not rely on any external libraries, e.g., for performing linear algebra routines. More specifically, all matrix factorizations and matrix-matrix multiplications are performed by standard triple loop implementations. It is well known that such textbook implementations can easily be outperformed, for medium- to large-scale matrix dimensions, by more advanced Basic Linear Algebra Subprogram (BLAS) libraries. For this purpose, high-performance open-source BLAS implementations could be used instead such as `GoToBLAS`⁴² and `OpenBLAS`⁴³, as well as proprietary implementations such as Intel's `MKL` and AMD's `ACML`. The open-source `BLASFEO`⁴⁴ framework provides hardware-tailored optimized dense BLAS routines, outperforming most alternative tools for small- to medium-scale matrix dimensions that occur typically for applications of embedded optimization. However, in what follows, we illustrate that our simplified self-contained C code implementation remains very competitive with state of the art QP solvers. The main reasons for this are that PRESAS exploits the block-structured sparsity in direct optimal control problems and it does not require any matrix factorizations or matrix-matrix multiplications within the active-set iterations.

In addition, the C code implementation is written such that the solver can easily be embedded in prototyping platforms such as, e.g., the `dSPACE MicroAutoBox-II` unit, as well as in microcontrollers that are used in real-world mechatronic systems. The computational hardware that we target for implementation has capabilities in the same scale of current automotive microcontrollers. As it is well known, due to the need to operate in harsh environmental conditions, the hard real-time requirement, and the cost considerations, current automotive microcontrollers are significantly less computationally powerful than standard desktop computers⁴, in terms of both memory and operations per second. This limits the amount of computations and data storage that we use in implementing our numerical optimization algorithms. Table 1 illustrates the variants of the proposed primal active-set

Algorithm 4 Preconditioned RESidual method within primal Active-Set (PRESAS) solver for MPC.

Input: Primal feasible starting point y_0 and working set \mathcal{W}_0 (see warm starting in Section 5).

- 1: Compute regularized Hessian block matrix $\tilde{H}_k \approx H_k$ such that $\tilde{H}_k > 0$ for $k = 0, \dots, N$.
- 2: **if** \tilde{H}_k is diagonal **then**
- 3: Compute \tilde{H}_k^{-1} by inverting diagonal elements for each $k = 0, \dots, N$. $\mathcal{O}(Nm)$
- 4: **else**
- 5: Compute Cholesky factorization of block matrix \tilde{H}_k for $k = 0, \dots, N$. $\mathcal{O}(Nm^3)$
- 6: Block-tridiagonal (reverse) Cholesky factorization for matrix $\mathcal{A}\tilde{H}^{-1}\mathcal{A}^\top$ in Eq. (15). $\mathcal{O}(Nm^3)$
- 7: **for** iter = 1, 2, ..., max **do**
- 8: Algorithm 2 to solve the linear system in (4), to compute the search direction $(\Delta y, \Delta v)$. $\mathcal{O}(Nm^2)$
- 9: $\alpha \leftarrow \min_{(k,j) \notin \mathcal{W}} -\frac{d_{k,j} + D_{k,j}\tilde{w}_k}{D_{k,j}\Delta w_k}$. (step length computation)
- 10: **if** $\alpha < 1$ **then**
- 11: $(k, j) \leftarrow \arg \min_{(k,j) \notin \mathcal{W}} -\frac{d_{k,j} + D_{k,j}\tilde{w}_k}{D_{k,j}\Delta w_k}$. (most blocking constraint)
- 12: $\mathcal{W} \leftarrow \mathcal{W} \cup \{(k, j)\}$. (add constraint to working set)
- 13: Add row to matrix \mathcal{A} and compute rank-one factorization update for $\mathcal{A}\tilde{H}^{-1}\mathcal{A}^\top$. $\mathcal{O}(Nm^2)$
- 14: $y^* \leftarrow y + \alpha \Delta y$ and $v^* \leftarrow v + \alpha \Delta v$.
- 15: **else**
- 16: $y^* \leftarrow y + \Delta y$ and $v^* \leftarrow v + \Delta v$.
- 17: **if** $\exists (k, j) \in \mathcal{W} : \mu_{k,j}^* < 0$ **then**
- 18: $\mathcal{W} \leftarrow \mathcal{W} \setminus \{\arg \min_{(k,j) \in \mathcal{W}} \mu_{k,j}^*\}$. (remove from working set)
- 19: Remove row from \mathcal{A} and compute rank-one factorization update for $\mathcal{A}\tilde{H}^{-1}\mathcal{A}^\top$. $\mathcal{O}(Nm^2)$
- 20: **else**
- 21: **stop** (QP solution found)

Output: Optimal solution y^* and v^* for the QP in (1).

TABLE 1 Most important variants of the primal active-set solver that have been implemented in self-contained C code.

	Sparsity structure	Preconditioned iterative solver	Primal-feasible warm starting
PRESAS-SC	Block-sparse	PMINRES - Schur complement (SC)	Forward simulation + slack variables
PRESAS-AL	Block-sparse	PMINRES - Augmented Lagrangian (AL)	Forward simulation + slack variables
PRESAS-PPCG	Block-sparse	PPCG - Constraint preconditioner (Alg. 2)	Forward simulation + slack variables
PRESAS-PPCG-WSA	Block-sparse	PPCG - Constraint preconditioner (Alg. 2)	Augmented Lagrangian method (Alg. 3)
D-PRESAS-PPCG	Dense	PPCG - Constraint preconditioner (Alg. 2)	Slack variables

method that are implemented in self-contained C code. Our implementation allows the flexibility of switching between different structure exploiting preconditioning techniques within the proposed block-sparse PRESAS solver, in order to adapt to a particular problem formulation. Table 1 also includes an implementation of our solver that is based on dense linear algebra, called D-PRESAS, which will be used in comparisons of the numerical simulation results in Section 7. By comparing D-PRESAS with PRESAS, based on the same preconditioned iterative solver, we can show the scalability of the tailored block-sparse structure exploitation for fast MPC implementations. As expected, D-PRESAS can remain competitive for relatively small-scale applications or it can even outperform the sparse PRESAS solver, and showing its capabilities allows us to assess the performance of the preconditioned iterative solver with respect to direct linear algebra methods.

6.2 | Nonlinear MPC: ACADO Code Generation Tool Interface

In addition to linear and linear time-varying MPC applications, we illustrate the computational performance of the PRESAS solver also for nonlinear MPC (NMPC) applications. For this purpose, we use a minimal software interface between the code generation tool of the ACADO Toolkit^{45,46} and the self-contained C code implementations of PRESAS. The nonlinear optimal control solver in the open-source toolkit uses an online variant of a sequential quadratic programming (SQP) type method,

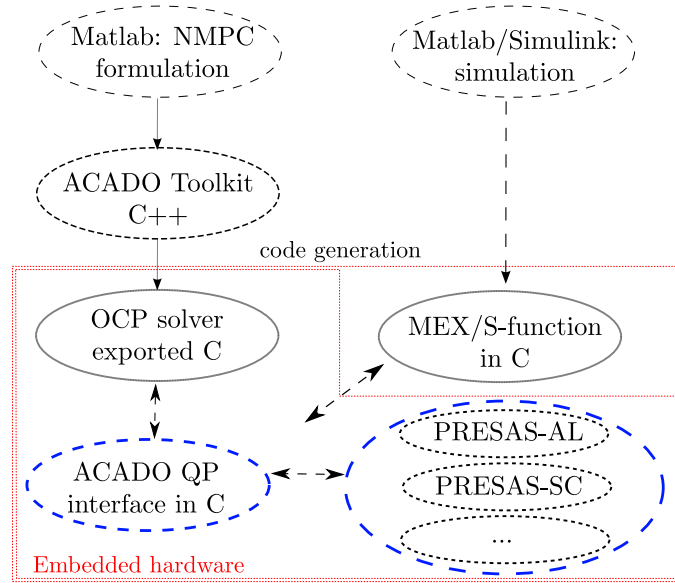


FIGURE 1 Illustration of the proposed work flow based on the open-source ACADO Toolkit software and its C code generation tool for nonlinear MPC: custom sparse QP solver interface for PRESAS.

known as the real-time iterations (RTI) scheme². The RTI approach is based on performing one SQP iteration per control time step, and on using a continuation-based warm starting of the state and control trajectories from one time step to the next. Each iteration consists of two steps:

1. *Preparation phase*: discretize and linearize the system dynamics, linearize the remaining constraint functions, and evaluate the quadratic objective approximation to build the optimal control structured QP subproblem of the form in (1).
2. *Feedback phase*: solve the local QP approximation (1) to update the current values for all optimization variables, i.e., the state and control trajectories, and obtain the next control input to apply to the system.

Note that we developed two separate interfaces for either the block-sparse and dense versions of the PRESAS solver in Table 1. The dense D-PRESAS variant can be used together with condensing²⁶, to numerically eliminate the state variables from the QP subproblem at each control time step, in the auto generated RTI optimization algorithm. The resulting work flow, based on our custom ACADO-PRESAS software interface, is illustrated in Figure 1 and consists of the following three steps:

1. *Formulation*: high-level problem formulation in Matlab or C++ using ACADO syntax, including the continuous-time nonlinear system dynamics, least squares type objective function and inequality constraints.
2. *Generation*: automatic generation of C code to implement an RTI algorithm for the user-specified optimal control problem. The generated code defines routines for a generalized Gauss-Newton (GGN) method that performs tailored function evaluations and algorithmic differentiation (AD). More specifically, the OCP solver performs numerical simulation and sensitivity propagation for discretization and linearization of the nonlinear system of differential equations, and it computes a linearization of the nonlinear inequality constraints and a Gauss-Newton type Hessian approximation.
3. *Compilation*: the generated ACADO code can be compiled together with the self-contained PRESAS solver and its custom interface, and linked into a MEX file or S-function such that it can be used from Matlab or Simulink.

7 | NUMERICAL SIMULATION CASE STUDIES

We consider three different numerical simulation case studies. The first case study concerns an MPC problem formulation based on linear system dynamics for a chain of spring-connected masses as in²⁹. The second case study is a nonlinear dynamic system that describes an inverted pendulum on top of a cart. We consider the case of linear MPC for stabilization of the inverted

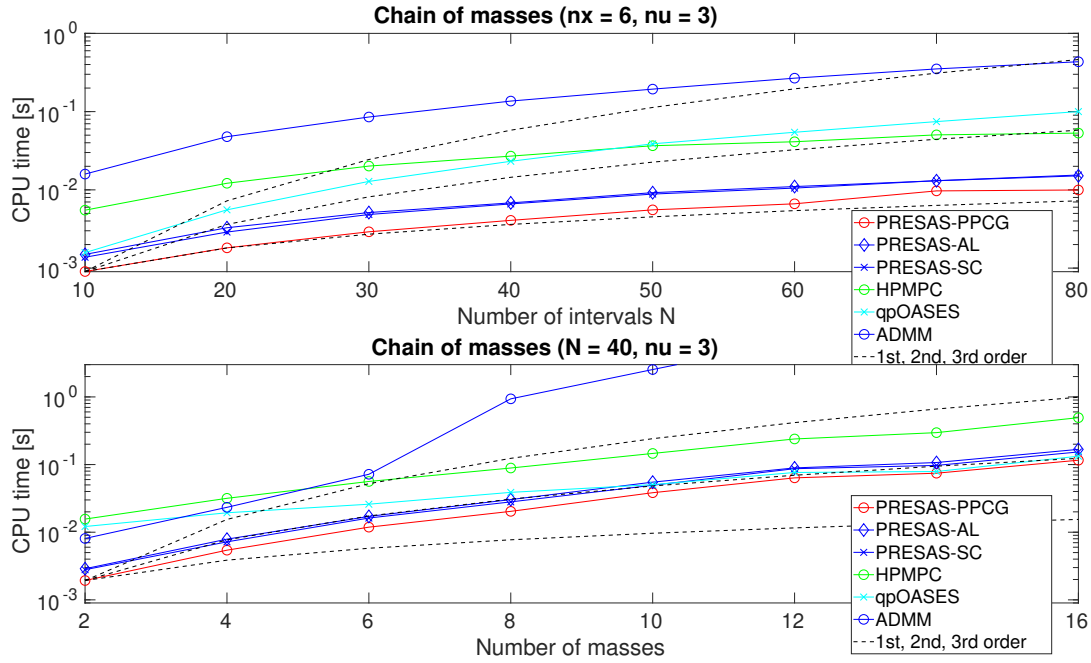


FIGURE 2 Average computation time (ARM Cortex-A7 processor in the Raspberry Pi 2) of different QP solvers for linear MPC on the chain of masses with L2 slack reformulation: varying number of control intervals N and number of masses n_m .

pendulum and nonlinear MPC for the swing-up of the pendulum⁴⁶. Finally, a hardware-in-the-loop simulation of a trajectory tracking NMPC algorithm for vehicle control is presented in the third case study. We illustrate the computational performance for different variants of the PRESAS solver in Table 1, based on standalone C code implementations.

7.1 | Case Study 1: MPC on Chain of Oscillating Masses

This first test problem consists of the chain of oscillating masses, which is often used as a benchmark example for fast MPC solvers^{20,29}. The linear time-invariant system dynamics and corresponding OCP formulation, of the form in (1), are described in more detail in²⁹. The full state of the system consists of the displacement and velocity of the n_m masses, i.e., $x(t) \in \mathbb{R}^{2n_m}$ such that the state dimension can be changed by changing the amount of masses. A number of actuators $n_u < n_m$ apply tensions between certain masses while respecting the actuator limitations as well as constraints on the position and velocity of each of the masses. In order to guarantee that the QP remains feasible at each sampling instant, the state constraints are softened by adding one slack variable for each control interval. Based on a sufficiently large penalization of this additional variable in the objective, a feasible solution can be found whenever possible. As discussed in⁴⁷, the ADMM solver does not require the explicit use of slack variables in the OCP problem formulation to soften state-dependent inequality constraints based on soft thresholding. Therefore, unlike for all the other optimization algorithms, the numerical results in this paper are obtained without including auxiliary slack variables for ADMM. During the closed-loop MPC simulations, a disturbance force, randomly generated with uniform distribution and kept the same for all simulations, acts on each of the masses as in²⁹.

Scaling Computational Complexity of PRESAS Solver with Problem Dimensions

Figure 2 shows the average computation time per QP solution during the closed-loop simulations of linear MPC for different numbers of masses n_m and for a varying control horizon length N . For simplicity, we do not consider the case where the control horizon is shorter than the prediction horizon, resulting in relatively less control variables, which could be used to reduce the computational cost of any optimization algorithm. Figure 2 includes the computation time for ADMM, qpOASES, HPMPC and the 3 variants of PRESAS. The considered ADMM algorithm²⁴ and qpOASES⁶ both solve the small but dense QP after numerically eliminating the state variables. The timing results for ADMM and qpOASES include the computation time for this condensing routine²⁶. Note that the black dashed lines in Figure 2 illustrate, respectively, a computational cost that increases with a 1st, 2nd

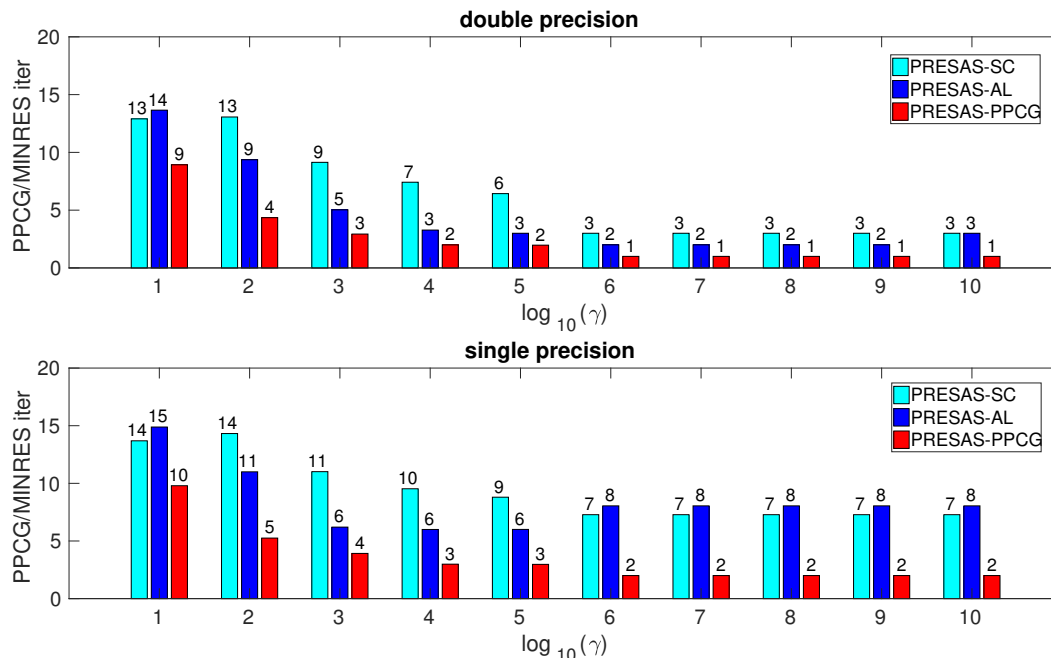


FIGURE 3 Average number of residual iterations for PPCG in PRESAS-PPCG or for PMINRES in PRESAS-SC or PRESAS-AL per linear system solution within the closed-loop MPC simulations for a chain of $n_m = 3$ masses.

or 3rd order of complexity, i.e., $\mathcal{O}(N)$, $\mathcal{O}(N^2)$ and $\mathcal{O}(N^3)$ in the top part, and $\mathcal{O}(m)$, $\mathcal{O}(m^2)$ and $\mathcal{O}(m^3)$ in the bottom part of Figure 2. The typical runtime complexity of $\mathcal{O}(N^2 m^2)$ for the two dense solvers, ADMM and qpOASES, can be observed. However, note that the computation time for qpOASES remains nearly constant with respect to the number of masses, because the dense QP includes only the control variables and Figure 2 is constructed for a fixed amount of $n_u = 3$ control inputs. When increasing the number of masses in the bottom part of Figure 2, both the ratio of control to state variables and the ratio of active inequality constraints decrease. On the other hand, the computational complexity of $\mathcal{O}(N m^3)$ for the interior-point method in HPMPC⁹ and the per iteration complexity of $\mathcal{O}(N m^2)$ for the proposed variants of the PRESAS solver can be observed in Figure 2. Even though the PRESAS-AL and PRESAS-SC solvers are very competitive with state of the art optimal control algorithms, it can be observed that the PRESAS-PPCG method further outperforms these algorithms. This performance scales well with the number of control intervals and the number of state variables in this particular case study.

Solver Sensitivity to Parameter Selection and Numerical Accuracy

PRESAS-AL requires a suitable choice for its parameter γ , while PRESAS-SC and the PRESAS-PPCG solver typically need diagonal regularization using a parameter $\epsilon = \frac{1}{\gamma}$. Figure 3 illustrates the numerical dependency of the number of PMINRES and PPCG iterations on the value of γ , either using single- or double-precision arithmetics. As expected, the value for γ should be sufficiently large but not too large in order to maintain a good conditioning of the preconditioner. In case of double-precision arithmetics, the parameter design appears to be much less sensitive. For a sufficiently large value of γ , the SC, AL and PPCG variant of PRESAS, respectively, require 3, 2 and 1 iteration of the residual method on average for each solution of a saddle point linear system. When using single-precision arithmetics, the numerical performance is reduced for all three preconditioners. The PRESAS-PPCG method shows a consistently lower number of iterations and it appears to be less sensitive to the choice of parameter value γ for the considered range of values. In addition, the PRESAS-PPCG method allows an average of only two PPCG iterations for a sufficiently small regularization parameter $\epsilon = \frac{1}{\gamma}$ in this case study. In order to avoid ill conditioning of the preconditioner, a good choice for the parameter value is, e.g., $\gamma = 10^7$ for double precision and $\gamma = 10^4$ for single-precision arithmetics. Note that this parameter value is fixed for the numerical results that are presented in the remainder of this article.




FIGURE 4 Closed-loop trajectories for linear MPC to stabilize an inverted pendulum on top of a cart, starting from different initial conditions for the pendulum angle ($\theta_0 = 0.04$, $\theta_0 = 0.12$ or $\theta_0 = 0.20$ rad).

7.2 | Case Study 2: MPC for Inverted Pendulum on a Cart

This second numerical case study consists of two closed-loop MPC simulations that, respectively, consider stabilization and swing-up of the inverted pendulum mounted on top of a cart⁴⁶. We present detailed timing results for an ARM Cortex-A7 processor in the Raspberry Pi 2. While they are not embedded processors by themselves, such Raspberry Pis use ARM cores with computational capabilities on the order of high-end microcontrollers that can be used for embedded real-time control applications in several industries⁴.

Warm-started Initialization for PRESAS solver: Inverted Pendulum Stabilization

This first simulation involves stabilizing the nonlinear inverted pendulum in the upward unstable position, based on a linearization of the nonlinear dynamics in that steady state. The closed-loop trajectories are illustrated in Figure 4 for different initial conditions for the angle of the pendulum $\theta_0 = 0.04$, 0.12 or 0.20 rad. Both the actuated force and the cart position are constrained to remain within their respective bounds. As illustrated in Figure 4, it becomes impossible to satisfy the position constraint, resulting in a constraint violation around $t = 1$ s in case of the initial value $\theta_0 = 0.20$, which is treated by a softened reformulation of the state-dependent inequality constraints. This results in a relatively high number of online active-set changes with respect to the amount of state $n_x = 4$ and control variables $n_u = 2$.

Table 2 shows the average and worst-case computation times on an ARM Cortex-A7 processor and the numbers of iterations for linear MPC of the inverted pendulum, including results for the QP solvers ADMM, qpOASES, HPMPC, qpDUNES and different variants of our proposed PRESAS solver. For all the solvers, except for the interior point method in HPMPC, warm starting is performed. For most of the QP solvers, the computational effort grows significantly when increasing the initial value for the angle of the inverted pendulum. The results show that the PRESAS-PPCG algorithm consistently outperforms both the PRESAS-SC and the PRESAS-AL variants of our solver. In combination with the augmented Lagrangian based warm-started initialization procedure, see Section 5.2, the resulting PRESAS-PPCG-WSA solver demonstrates an improved worst-case computational performance. Note that qpOASES reports 0 iterations in Table 2 when the initial guess for the active set is optimal. And the results for the PRESAS-PPCG-WSA solver refer to the total number of inner iterations in Algorithm 3, i.e., it corresponds to the total number of active-set changes, which is equal to the total number of linear systems that are solved to compute each QP solution. More specifically, for this particular case study, our solver is at least a factor 2 times faster than all other state of the art QP solvers in

TABLE 2 Average and worst-case computation times (ms) for MPC of an inverted pendulum ($\mathcal{T}_s = 50$ ms, $N = 50$), with varying initial conditions for the angle of the pendulum, using double precision on an ARM Cortex-A7 processor in the Raspberry Pi 2.¹

	$\theta_0 = 0.04$		$\theta_0 = 0.12$		$\theta_0 = 0.20$	
	time [ms] (mean/max)	# iter (mean/max)	time [ms] (mean/max)	# iter (mean/max)	time [ms] (mean/max)	# iter (mean/max)
ADMM	8.10/63.56	17.9/92	17.24/90.51	42.1/181	25.94/109.79	65.4/239
qpOASES	7.25/16.24	0.0/0	7.71/28.26	0.1/3	20.84/67.71	2.8/10
HPMPC	15.80/17.38	9.0/10	17.46/21.84	9.5/11	18.95/24.43	10.4/13
qpDUNES	3.11/10.02	2.0/2	3.72/17.60	2.1/4	5.24/26.90	2.5/7
PRESAS-SC	2.30/6.39	1.1/4	2.65/23.59	1.2/16	4.43/43.42	2.6/30
PRESAS-AL	2.51/7.48	1.1/4	2.75/24.50	1.2/16	4.87/46.28	2.6/30
PRESAS-PPCG	1.38/2.91	1.1/4	1.58/9.46	1.2/16	2.50/17.73	2.6/30
PRESAS-PPCG-WSA	1.38/2.24	1.0/2	1.58/4.89	1.1/6	1.81/11.41	1.5/19

Table 2. In addition, by using the primal active-set method in PRESAS, a feasible but suboptimal solution can always be obtained by limiting the maximum number of active-set changes.

Nonlinear MPC Simulation Results: Inverted Pendulum Swing-up

The second numerical simulation involves a swing-up of a pendulum to its upward unstable position, using the nonlinear system dynamics and the optimal control problem formulation from⁴⁶. Both the actuated force and the cart position are constrained to remain within their respective bound values. A softened reformulation of the position constraints is used to guarantee each QP subproblem to remain feasible. The swing-up maneuver results in a relatively high number of online active-set changes with respect to the amount of state $n_x = 4$ and control variables $n_u = 2$. The nonlinear MPC (NMPC) algorithm is implemented using the Gauss-Newton based RTI scheme in the ACADO code generation tool as discussed in Section 6.2.

Table 3 shows the average and worst-case computation times for each QP solution within the RTI algorithm for NMPC of the inverted pendulum, including the solvers ADMM, qpOASES and three variants of PRESAS based on single- versus double-precision arithmetics on an ARM Cortex-A7 processor. The solver tolerance for ADMM was chosen to be equal to 10^{-3} , because first-order methods are generally not efficient in computing high-accuracy solutions⁴⁷. Instead, a solver tolerance of 10^{-6} was chosen for PRESAS and the accuracy of the solutions is relatively similar to that from qpOASES. When switching from double- to single-precision arithmetics, there can be an increased efficiency of the floating-point operations because, e.g., the compiler can more effectively use SIMD instructions. For example, this is observed for the computation time of the ADMM solver. However, the use of single-precision arithmetics may additionally lead to numerical issues that are caused by round-off errors, resulting in an overall increased number of iterations. One can observe from Table 3 that these competing effects result in a comparable computational performance of the proposed PRESAS-PPCG solver for both single- and double-precision arithmetics, which outperforms the alternative optimal control algorithms on this particular case study. On the other hand, PRESAS-SC and AL appear to be more negatively affected by the reduced precision, while remaining competitive with the other solvers.

7.3 | Case Study 3: Hardware-in-the-Loop Simulations of NMPC for Vehicle Control

The third case study illustrates the real-time computational feasibility and the memory footprint of our proposed PRESAS solver for NMPC based trajectory tracking in an autonomous driving application. The system dynamics are based on a single-track vehicle model, obtained by lumping together the left and right wheel on each axle, coupled with front and rear wheel speed dynamics and torque generation and transfer dynamics and first order dynamics for the front steering actuator. The tire friction forces are described by Pacejka's magic formula, including first order dynamics for the front and rear slip angles, and the coupling between longitudinal and lateral tire forces is based on the friction ellipse. The trajectory tracking objective is formulated by introducing a time-dependent polynomial for each of the outputs to be tracked as an approximation of the reference motion plan. Tracking a time-dependent plan can sometimes lead to large tracking errors even when the resulting motion closely corresponds to the planned motion, e.g., caused by the vehicle slowing down or speeding up relative to the reference. Instead, the reference

¹The Raspberry Pi 2 uses a BCM2836 SoC with a 900 MHz 32-bit quad-core ARM Cortex-A7 processor, with 256 KB shared L2 cache.

TABLE 3 NMPC for inverted pendulum swing-up ($T_s = 50$ ms and $N = 40$): total computation time and # solver iterations per call, # residual iterations per active-set change, using double- and single-precision arithmetics on an ARM Cortex-A7 processor.

Double precision	Time [ms] (mean/max)	# sol iter (mean/max)	# res iter (mean/max)	Single precision	Time [ms] (mean/max)	# sol iter (mean/max)	# res iter (mean/max)
PRESAS-PPCG	1.82/6.90	2.5/12	1.0/1	PRESAS-PPCG	1.81/7.37	2.5/12	1.9/2
PRESAS-AL	3.28/14.59	2.5/12	2.0/3	PRESAS-AL	4.75/22.19	2.5/12	6.4/8
PRESAS-SC	3.67/15.81	2.5/12	3.0/3	PRESAS-SC	6.85/24.67	5.4/21	5.0/6
qpOASES	16.96/34.86	4.2/14	-	qpOASES	16.37/38.86	5.2/16	-
ADMM	15.87/97.57	54.2/446	-	ADMM	11.68/69.61	54.5/446	-

trajectories are parameterized with respect to a path variable, that is defined by an additional state equation and for which the change rate can be directly controlled by the NMPC controller. This results in an additional degree of freedom, such that NMPC can accelerate/decelerate along the path. The inequality constraints in the NMPC problem formulation consist of geometric and physical limitations of the system, such as constraints on the distance of the vehicle position to the reference trajectory and on the lateral acceleration. These state-dependent constraints are formulated as soft inequality constraints, including an L1 penalty on the slack variable in the objective function. In addition, hard bounds are imposed on the steering wheel angle, steering rate, and on the wheel torques. The resulting OCP in the NMPC controller includes $n_x = 15$ differential states, $n_u = 5$ control inputs and $N = 20$ or $N = 30$ control intervals for a $T = 2$ or $T = 3$ s horizon length, respectively. The vehicle and tire friction model parameters are from⁴⁸. More details on the problem formulation can be found in^{49,50}.

We evaluate the controller in a dSPACE MicroAutoBox-II rapid prototyping unit (RPU), which is equipped with a 900 MHz PowerPC real-time processor (IBM PPC 750GL) and 16 MB of RAM. The capabilities of the RPU are slightly superior, but relatively on the same order, compared to those of current and next-generation processors for hard real-time automotive applications. We assess the real-time computational feasibility of the proposed NMPC controller based on an automatically generated C code implementation of the RTI algorithm, using the ACADO code generation tool, in combination with different embedded QP solvers. Table 4 shows the average and worst-case closed-loop computation times using a sampling time of $T_s = 100$ ms and two different values for the number of control intervals N_{MPC} on the dSPACE MicroAutoBox-II. In addition, Table 4 presents the total memory usage for each of the software implementations of the NMPC controller. The QP solution time scales linearly with respect to the horizon length for the proposed PRESAS solvers and the total computation time remains well below the hard real-time sampling requirement of 100 ms. Note that all three PRESAS variants in Table 4 are based on the PPCG solver. The results based on dense QP solvers, i.e., D-PRESAS, qpOASES and ADMM, are obtained using the dense QP interface in ACADO that includes a condensing routine to numerically eliminate the state variables from the sparse QP subproblem at each control time step. As discussed in⁴⁷, the ADMM solver does not require the explicit use of slack variables in the OCP problem formulation to soften state-dependent inequality constraints, resulting in a reduced preparation time compared to D-PRESAS and qpOASES in Table 4. Table 4 shows the percentage of overruns, i.e., the time steps in which the total computation time exceeds the desired sampling time of 100 ms. The ADMM solver occasionally does not reach the required solver tolerance within $T_s = 100$ ms in case $N_{\text{MPC}} = 20$ and the overrun percentage reaches even a value of 42 % in case $N_{\text{MPC}} = 30$. The state of the art qpOASES solver is quite competitive with the different variants of our PRESAS solver for this NMPC simulation case study if $N_{\text{MPC}} = 20$. However, for a longer control horizon, the total memory usage of the software implementation exceeds the amount of memory that is available on the dSPACE MicroAutoBox-II. On the other hand, the dense linear algebra variant of our QP solver, D-PRESAS, is more compact than qpOASES, such that this NMPC implementation fits in the dSPACE MicroAutoBox-II even for $N_{\text{MPC}} = 30$.

It can be seen from Table 4 that all three variants of the PRESAS solver are real-time feasible and result in a relatively low memory footprint. The sparse implementation of the PRESAS-WSA method results overall in the best computational performance, especially when looking at the worst-case computation time for a larger problem size, i.e., $N_{\text{MPC}} = 30$ in Table 4. More specifically, one can observe that the maximum number of iterations increased from 9 to 11 for PRESAS-WSA, instead of the increase from 8 to 26 for PRESAS, when increasing the control horizon length from $N_{\text{MPC}} = 20$ to $N_{\text{MPC}} = 30$. The latter is a good motivation for the augmented Lagrangian type initialization in PRESAS-WSA (see Section 5.2) instead of the forward simulation based initialization procedure in Section 5.1 for the PRESAS solver. Forward simulation of unstable vehicle dynamics can lead to a bad initial guess for the primal active-set strategy, which is an issue that becomes more obvious for an increasingly long prediction horizon. For a smaller problem size, i.e., $N_{\text{MPC}} = 20$, the dense implementation (D-PRESAS) is still competitive, and

its comparison with other dense solvers provides an assessment of the performance of the primal active-set strategy with the preconditioned iterative solver on its own (without the block-structured sparsity exploitation).

$N_{\text{MPC}} = 20, T_s = 100$ ms	PRESAS-WSA	PRESAS	D-PRESAS	qpOASES	ADMM
NMPC – QP preparation time	7.3/7.4 ms	7.3/7.4 ms	12.7/12.8 ms	12.7/12.8 ms	11.0/11.0 ms
NMPC – QP solution time	12.7/22.2 ms	11.4/20.6 ms	11.3/18.0 ms	10.7/25.2 ms	43.1/137.1 ms
NMPC – QP solver: # iter	3/9	2/8	1/5	2/23	107/443
Total time (mean/max)	20.0/29.6 ms	18.7/28.0 ms	24.0/30.8 ms	23.4/38.0 ms	54.1/148.1 ms
Total overrun percentage	0 %	0 %	0 %	0 %	7.4 %
Total memory usage	8.56 MB	8.54 MB	8.88 MB	15.13 MB	8.46 MB
$N_{\text{MPC}} = 30, T_s = 100$ ms	PRESAS-WSA	PRESAS	D-PRESAS	qpOASES	ADMM
NMPC – QP preparation time	11.2/11.2 ms	11.2/11.3 ms	22.3/22.4 ms	–	18.7/18.8 ms
NMPC – QP solution time	19.0/34.6 ms	18.5/71.3 ms	31.5/52.4 ms	–	75.2/197.6 ms
NMPC – QP solver: # iter	3/11	2/26	1/6	–	51/249
Total time (mean/max)	30.2/45.8 ms	29.7/82.6 ms	53.8/74.8 ms	–	93.9/216.4 ms
Total overrun percentage	0 %	0 %	0 %	–	42.4 %
Total memory usage	9.04 MB	9.02 MB	10.1 MB	–	9.12 MB

TABLE 4 Average and worst-case timing results (ms) for a 300 s simulation of the NMPC based vehicle control algorithm using a sampling time $T_s = 100$ ms and $N_{\text{MPC}} = 20$ or $N_{\text{MPC}} = 30$ on dSPACE MicroAutoBox-II. All three PRESAS variants are based on the PCG solver and the results for the dense QP solvers, D-PRESAS, qpOASES and ADMM, are obtained using condensing.

8 | EXPERIMENTAL RESULTS FOR NMPC OF AUTONOMOUS SCALED VEHICLES

In this section, we present experimental results for an NMPC implementation, using the proposed PRESAS solver, of a real-time trajectory tracking controller for an autonomous driving system, implemented on a testbench of small-scale vehicles^{51,52}. Such a platform, based on the Hamster robots as illustrated in Fig. 5a, can be used for rapid prototyping and evaluation of a complete control and estimation software stack before deploying to full-scale vehicle experiments, e.g., see^{53,54}. Next, we briefly describe the test setup and the NMPC problem formulation, before presenting the experimental results in Section 8.3.

8.1 | Scaled Vehicle Experimental Platform

The Hamster is a 25×20 cm mobile robot for research and prototype development, see Fig. 5a. It is equipped with scaled versions of sensors commonly available on full-scale research vehicles, such as a 6 m range mechanically rotating 360 deg Lidar, an inertial measurement unit (IMU), GPS receiver, HD camera, and motor encoders. It uses two Raspberry Pi 3 computing platforms, each with an ARM Cortex-A53 processor. The robot has Ackermann steering and therefore it is kinematically equivalent to a full-scale vehicle, and its dynamic behavior resembles that of a regular vehicle. The Hamster has low-level dedicated hardware for power distribution and monitoring. At the lowest layer, the robot is controlled by setting the desired front wheel steering angle and longitudinal velocity. In addition, the Hamster has built-in mapping and localization capabilities, and object detection and tracking can be done with the onboard Lidar and/or camera. Hence, the platform is a good test setup for verifying the dynamic feasibility and performance of vehicle control and estimation software in a realistic setting, with a sensor setup similar to the one expected in full-scale autonomous vehicles. The Hamster communicates and connects to external algorithms using the robot operating system (ROS). In order to be able to accurately evaluate the control and estimation algorithms in terms of performance in a controlled environment, we use an OptiTrack motion capture system. The OptiTrack system is a flexible

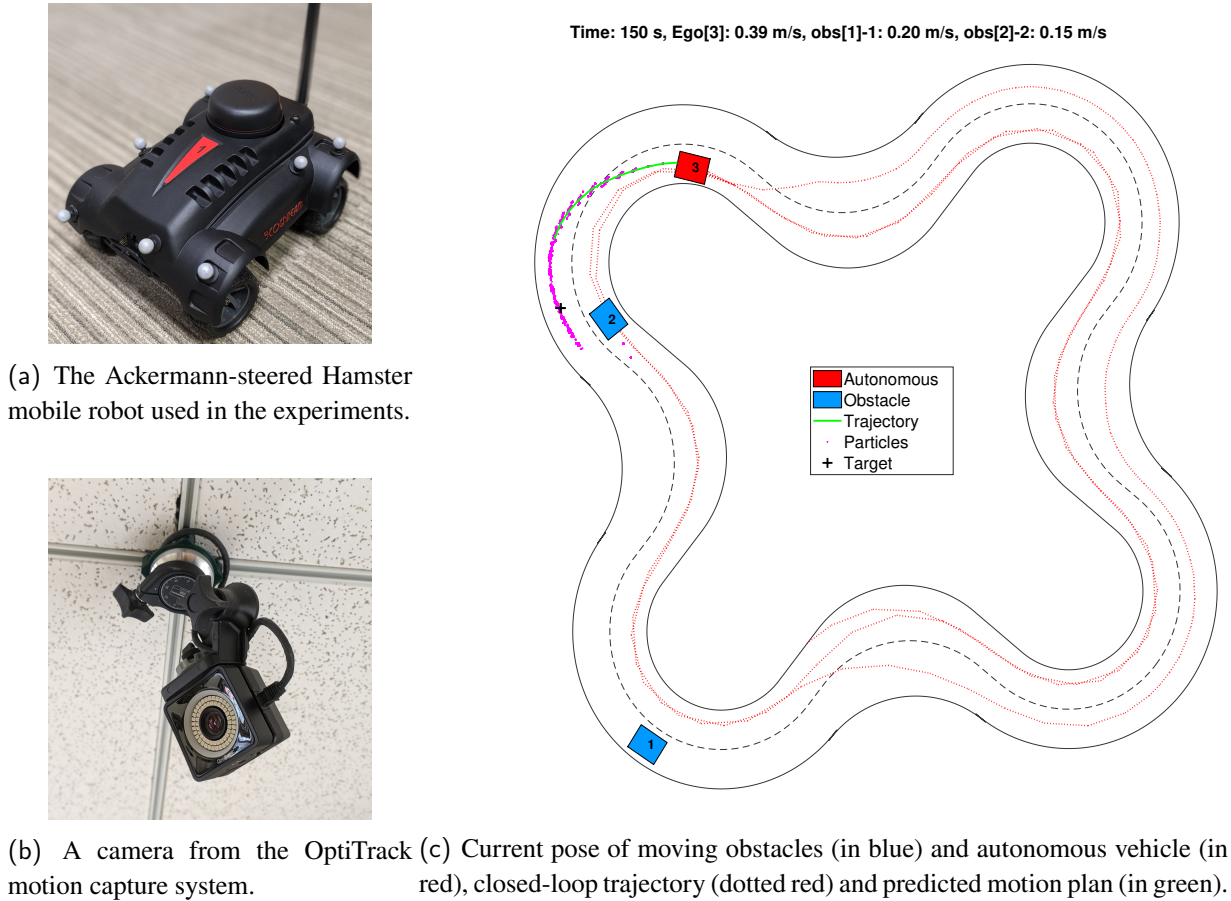


FIGURE 5 Illustration of our experimental testbench that consists of small-scale test vehicles (a), with on-board sensors, and an OptiTrack motion capture system (b). Figure (c) shows the track and a snapshot of the position of the autonomous vehicle and of the two obstacle vehicles, as well as an illustration of the closed-loop trajectory at the end of the experiment.

camera-based (see Fig. 5b) six degrees-of-freedom tracking system that can be used for tracking drones, ground vehicles, and industrial robots. Depending on the environment and quality of the calibration, the system can track the position of the Hamster robot within 0.9 mm accuracy and with a rotational error of less than 3 deg. The OptiTrack is connected to the same ROS network using the common VRPN protocol.

8.2 | Vehicle Dynamics and NMPC Formulation

Due to the testing environment, we consider a kinematic single-track model in which the two wheels on each axle are lumped together. Although a dynamic vehicle model based on force balances is generally more accurate than a kinematic model, the latter may be still sufficient for several driving conditions⁵⁵, and model errors are corrected for by the feedback control. On the other hand, a dynamic model depends on more parameters, such as the wheel radii, tire stiffness, and vehicle mass and inertia, which typically are unknown/uncertain and may be difficult, or at least tedious, to estimate. Here, for simplicity, we use the

kinematic single-track model

$$\dot{x} = \begin{bmatrix} \dot{p}_X \\ \dot{p}_Y \\ \dot{\psi} \\ \dot{\delta}_f \\ \dot{v}_x \\ \dot{\delta} \end{bmatrix} = \begin{bmatrix} v_x \cos(\psi + \beta) \\ v_x \sin(\psi + \beta) \\ \frac{v_x}{L} \tan(\delta_f) \cos(\beta) \\ 1/t_d(\delta + \delta_0 - \delta_f) \\ u_1 \\ u_2 \end{bmatrix}, \quad (20)$$

where p_X, p_Y is the longitudinal and lateral position in the world frame, ψ is the heading angle and $\dot{\psi}$ the heading rate of the vehicle, v_x is the longitudinal velocity of the vehicle, δ and δ_f are, respectively, the desired and actual front wheel steering angle, $L := l_f + l_r$ is the wheel base, and $\beta := \arctan(l_r \tan(\delta_f)/L)$ is the body-slip angle. A first order front steering equation is included in Eq. (20) which models that, due to the steering mechanism and other platform effects, the wheel angle response is not immediate. In addition, an input additive offset value δ_0 can be estimated online. The inputs u_1, u_2 are the acceleration and the steering rate, respectively. This choice allows to provide smooth velocity and steering profiles and to constrain the allowed change rate of the vehicle velocity and front steering wheel angle.

Similar to the vehicle control hardware-in-the-loop simulation results in Section 7.3, the NMPC trajectory tracking objective is formulated based on a smooth approximation of the reference motion plan at each control time step, parameterized with respect to a time-dependent path variable τ for which the change rate $\frac{d\tau}{dt} = \dot{\tau}$ is an additional control input. In continuous time, the NMPC cost function therefore consists of the following terms

$$\int_0^T \left(\|F^{\text{ref}}(x(t)) - y^{\text{ref}}(\tau, d)\|_W^2 + \|x(t)\|_Q^2 + \|u(t)\|_R^2 + r_s s(t) \right) dt, \quad (21)$$

including a term for tracking the reference motion plan and two regularization terms for penalizing state and control variables, where $Q \in \mathbb{R}^{n_x \times n_x}$ and $R \in \mathbb{R}^{n_u \times n_u}$ are the corresponding weighting matrices. The definition of a positive slack variable $s(t) \geq 0$ allows for an easy implementation of an exact L1 soft constraint penalty in the objective, i.e., $|s(t)| = s(t)$. The reference motion plan is approximated by a smooth function $y^{\text{ref}}(\tau, d)$ that depends on the path variable τ as well as on additional parameters d . We define the tracking function in the objective based on a polynomial approximation of the reference trajectories as

$$F^{\text{ref}}(x(t)) - y^{\text{ref}}(\tau, d) = \begin{bmatrix} e_Y(t, \tau, d) \\ p_X(t) - p_X^{\text{ref}}(\tau, d_X) \\ p_Y(t) - p_Y^{\text{ref}}(\tau, d_Y) \\ \psi(t) - \psi^{\text{ref}}(\tau, d_\psi) \\ v_x(t) - v_x^{\text{ref}}(\tau, d_v) \end{bmatrix} = \begin{bmatrix} e_Y(t, \tau, d_X, d_Y, d_\psi) \\ p_X(t) - \sum_{i=0}^{n_X} d_X^i \tau^i \\ p_Y(t) - \sum_{i=0}^{n_Y} d_Y^i \tau^i \\ \psi(t) - \sum_{i=0}^{n_\psi} d_\psi^i \tau^i \\ v_x(t) - \sum_{i=0}^{n_v} d_v^i \tau^i \end{bmatrix}, \quad (22)$$

where the path error $e_Y(\cdot) = \cos(\psi^{\text{ref}}(\tau, d_\psi)) (p_Y - p_Y^{\text{ref}}(\tau, d_Y)) - \sin(\psi^{\text{ref}}(\tau, d_\psi)) (p_X - p_X^{\text{ref}}(\tau, d_X))$ is defined as the orthogonal distance to the parameterized reference trajectory. The inequality constraints in the NMPC problem formulation include hard bounds on the control inputs and soft constraints for limiting the distance to the parameterized reference trajectory, the vehicle velocity and the front steering wheel angle:

$$1 - \Delta \bar{\tau} \leq \dot{\tau} \leq 1 + \Delta \bar{\tau}, \quad -\bar{\delta} \leq \dot{\delta} \leq \bar{\delta}, \quad -\bar{v}_x \leq \dot{v}_x \leq \bar{v}_x, \quad 0 \leq s, \quad (23a)$$

$$-\bar{e}_Y \leq e_Y + s, \quad -\bar{\delta}_f \leq \delta_f + s, \quad -\bar{v}_x \leq v_x + s, \quad (23b)$$

$$e_Y - s \leq \bar{e}_Y, \quad \delta_f - s \leq \bar{\delta}_f, \quad v_x - s \leq \bar{v}_x. \quad (23c)$$

The resulting OCP includes $n_x = 7$ differential states, $n_u = 4$ control inputs and $N = 60$ control intervals with a sampling time of $T_s = 25$ ms over a $T = 1.5$ s horizon length. The NMPC controller is implemented with a sampling frequency of 40 Hz, using the direct multiple shooting discretization method in combination with the RTI scheme in the ACADO code generation tool and our proposed PRESAS optimization algorithm.

8.3 | Experimental Results and Solver Performance

We use three Hamsters in the experimental validation, one acts as the controlled ego vehicle (EV) and two Hamsters act as dynamic obstacles. The overall objective is to avoid the obstacles while circulating a two-lane closed circuit in the counter-clockwise direction, see Fig. 5c, with the inner lane as preferred lane. The obstacle vehicles (OVs) are commanded to track




FIGURE 6 Closed-loop results from Hamster experiments (see Figure 5): commanded front steering wheel angle δ , vehicle speed v_x , number of PRESAS solver iterations and computation time at each control time step.

the middle of either of the lanes, where the preferred lane and vehicle velocities can be changed throughout the experiments. The OV's use PID controllers for tracking the desired lane and velocity. The current position of each OV is obtained from the OptiTrack while its velocity is estimated. For the EV, the desired path trajectory, front steering angle, and vehicle velocity are computed by the motion planner and sent to the NMPC controller that tracks the trajectory in real time at a sampling frequency of 40 Hz. For this purpose, a particle filtering based motion planner is used, as presented recently in⁵². The planning algorithm typically runs at a relatively low frequency of 1-2 Hz and therefore relies on the NMPC controller to make the vehicle accurately execute the desired motion. The current state of the vehicle, as well as disturbances such as the steering offset value, are estimated online using an extended Kalman filter based on IMU and OptiTrack measurements.

Figure 6 shows closed-loop results from 150 s of Hamster experiments on the test track in Figure 5c. The upper plot shows time trajectories of the control inputs that were sent by the NMPC controller to the actuators of the EV. The target velocity for the EV was chosen to be 0.4 m/s. It can be observed that the EV needed to slow down to a velocity of about 0.15 m/s between 50 and 100 s, since both lanes were blocked by slow moving OV's during this time period. The lower plot of Figure 6 illustrates the computational performance of our proposed PRESAS-PPCG-WSA implementation in Algorithm 3. Due to the warm starting of the solver, the number of active-set changes remains typically low at each control time step. Occasionally, due to a bad initial guess, the number of solver iterations can jump up to a relatively high number. But even the worst-case performance results in a total computation time that is well below the desired sampling time of $T_s = 25$ ms in these experiments.

9 | CONCLUSIONS

This paper presented different variants of a tailored implementation of a primal active-set strategy for solving optimal control structured quadratic programming problems as they typically arise in both linear and nonlinear real-time MPC applications. The proposed PRESAS solver is based on either a sparse block-diagonal preconditioner for the preconditioned minimal residual (PMINRES) method or a block-sparse constraint preconditioner for the projected preconditioned conjugate gradient (PPCG) method within a primal active-set optimization algorithm. In addition, we described different warm-started initialization procedures to compute a primal feasible starting point in a real-time predictive controller. Based on an efficient and self-contained C code implementation, the computational performance and numerical reliability of PRESAS is illustrated against other state of

the art QP solvers. We also showed that the PRESAS QP solver is real-time feasible on a dSPACE MicroAutoBox-II rapid prototyping unit for vehicle control applications. Finally, we presented experimental results for a real-time NMPC implementation based on PRESAS on a testbench of small-scale autonomous vehicles.

References

1. Rawlings JB, Mayne DQ, Diehl MM. *Model Predictive Control: Theory, Computation, and Design*. Nob Hill. 2nd edition ed. 2017.
2. Diehl M, Bock HG, Schlöder JP. A Real-Time Iteration Scheme for Nonlinear Optimization in Optimal Feedback Control. *SIAM Journal on Control and Optimization* 2005; 43(5) : 1714–1736
3. 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* 2020; 93(1) : 62-80.
4. Di Cairano S, Kolmanovsky IV. Real-time optimization and model predictive control for aerospace and automotive applications. In: *Proc. American Control Conf.*; 2018: 2392–2409.
5. Ferreau HJ, Almer S, Verschueren R, et al. Embedded Optimization Methods for Industrial Automatic Control. In: *IFAC World Congr.*; 2017: 13194–13209.
6. Ferreau HJ, Kirches C, Potschka A, Bock HG, Diehl M. qpOASES: a parametric active-set algorithm for quadratic programming. *Math. Progr. Comp.* 2014; 6(4) : 327–363.
7. Andersen MS, Dahl J, Vandenberghe L. CVXOPT: A Python package for convex optimization, version 1.1.6. Available at cvxopt.org; 2013.
8. Domahidi A, Perez J. FORCES Professional. embotech GmbH (<http://embotech.com>); 2014.
9. Frison G, Sorensen HB, Dammann B, Jørgensen JB. High-performance small-scale solvers for linear Model Predictive Control. In: *Proc. European Control Conf. (ECC)*; 2014: 128–133.
10. Benzi M, Golub GH, Liesen J. Numerical solution of saddle point problems. *Acta Num.* 2005 : 1–137.
11. Greenbaum A. *Iterative Methods for Solving Linear Systems*. SIAM . 1997
12. Boland D, Constantinides GA. An FPGA-based implementation of the MINRES algorithm. In: *Inter. Conf. Field Prog. Logic and Appl.*; 2008: 379-384
13. Knyazev A, Fujii Y, Malyshev A. Preconditioned Continuation Model Predictive Control. In: *Proc. IEEE Conf. Control Appl. (CCA)*; 2015: 101–108.
14. Shahzad A, Kerrigan EC, Constantinides GA. A fast well-conditioned interior point method for predictive control. In: *49th IEEE Conference on Decision and Control (CDC)*; 2010: 508-513
15. Shahzad A, Kerrigan E, Constantinides G. Preconditioners for inexact interior point methods for predictive control. In: *American Control Conference (ACC), 2010*; 2010: 5714–5719.
16. Otta P, Burant J, Šantin O, Havlena V. Newton projection with proportioning using iterative linear algebra for model predictive control with long prediction horizon. *Optimization Methods and Software* 2019; 34(5) : 1075-1098.
17. Bartlett R, Wächter A, Biegler L. Active Set vs. Interior Point Strategies for Model Predictive Control. In: *Proc. ACC*; 2000; Chicago, Il: 4229–4233.
18. Wright S. Applying new optimization algorithms to model predictive control. In: Kantor J, Garcia C, Carnahan B., eds. *Fifth Intern. Conf. on Chemical Process Control – CPC V* 1996 (pp. 147–155).

19. Kirches C, Bock HG, Schlöder JP, Sager S. A factorization with update procedures for a KKT matrix arising in direct optimal control. *Math. Progr. Comp.* 2011; 3(4) : 319–348
20. Quirynen R, Knyazev A, Di Cairano S. Block Structured Preconditioning within an Active-Set Method for Real-Time Optimal Control. In: *Proc. European Control Conf. (ECC)*; 2018: 1154–1159.
21. Quirynen R, Knyazev A, Di Cairano S. Projected Preconditioning within a Block-Sparse Active-Set Method for MPC. In: *Proc. IFAC Conf. Nonlinear Model Predictive Control (NMPC)*; 2018: 28–34.
22. Nocedal J, Wright SJ. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering Springer. 2 ed. 2006.
23. Di Cairano S, Brand M, Bortoff SA. Projection-free parallel quadratic programming for linear model predictive control. *Inter. J. of Control* 2013; 86(8) : 1367–1385.
24. Raghunathan AU, Di Cairano S. ADMM for Convex Quadratic Programs: Q-Linear Convergence and Infeasibility Detection. *arXiv:1411.7288* 2015.
25. Bock HG, Plitt KJ. A Multiple Shooting Algorithm for Direct Solution of Optimal Control Problems. In: *Proc. IFAC World Congr.* Pergamon Press; 1984: 242–247.
26. Frison G, Jørgensen JB. A Fast Condensing Method for Solution of Linear-Quadratic Control Problems. In: *Proc. IEEE Conf. Decision and Control (CDC)*; 2013: 7715–7720.
27. Frasch JV, Sager S, Diehl M. A parallel quadratic programming method for dynamic optimization problems. *Math. Progr. Comp.* 2015; 7(3) : 289–329.
28. Anderson E, C. ZB, Bischof , et al. *LAPACK Users' Guide*. Philadelphia, PA: SIAM. third ed. 1999.
29. Wang Y, Boyd S. Fast Model Predictive Control Using Online Optimization. *IEEE Transactions on Control Systems Technology* 2010; 18(2) : 267–278.
30. Frison G, Jørgensen JB. Efficient implementation of the Riccati recursion for solving linear-quadratic control problems. In: *2013 IEEE Inter. Conf. Control Appl. (CCA)*; 2013: 1117-1122
31. Arnström D, Axehill D. Exact Complexity Certification of a Standard Primal Active-Set Method for Quadratic Programming. In: *Proc. IEEE Conf. Decision and Control (CDC)*; 2019: 4317-4324.
32. Arnström D, Axehill D. A Unifying Complexity Certification Framework for Active-Set Methods for Convex Quadratic Programming. *arXiv preprint arXiv:2003.07605* 2020.
33. Benzi M, Wathen AJ. *Some Preconditioning Techniques for Saddle Point Problems*: 195–211; Berlin, Heidelberg: Springer Berlin Heidelberg . 2008
34. Verschueren R, Zanon M, Quirynen R, Diehl M. A Sparsity Preserving Convexification Procedure for Indefinite Quadratic Programs arising in Direct Optimal Control. *SIAM J. of Optim.* 2017; 27(3) : 2085–2109.
35. Greif C, Schoetzau D. Preconditioners for Saddle Point Linear Systems with Highly Singular (1,1) Blocks. *Electr. Trans. on Num. Anal.* 2006; 22 : 114–121.
36. Murphy MF, Golub GH, Wathen AJ. A Note on Preconditioning for Indefinite Linear Systems. *SIAM J. Scientific Computing* 2000; 21(6) : 1969-1972
37. Coleman TF, Verma A. A Preconditioned Conjugate Gradient Approach to Linear Equality Constrained Minimization. *Comp. Opt. Appl.* 2001; 20(1) : 61–72
38. Gould NIM, Hribar ME, Nocedal J. On the Solution of Equality Constrained Quadratic Programming Problems Arising in Optimization. *SIAM J. Sci. Comput.* 2001; 23(4) : 1376–1395

39. Bakhvalov NS, Knyazev AV. Preconditioned Iterative Methods in a Subspace for Linear Algebraic Equations with Large Jumps in the Coefficients. *Contemporary Mathematics* 1994 : 157–162.
40. Dollar H. *Iterative Linear Algebra for Constrained Optimization*. PhD thesis. University of Oxford, 2005.
41. Fletcher R. *Practical Methods of Optimization*. Chichester: Wiley. 2nd ed. 1987.
42. Goto K, Geijn RAVd. Anatomy of High-performance Matrix Multiplication. *ACM Trans. Math. Softw.* 2008; 34(3) : 12:1–12:25.
43. OpenBLAS . OpenBLAS: An optimized BLAS library. <http://www.openblas.net/>; 2011.
44. Frison G, Kouzoupis D, Zanelli A, Diehl M. BLASFEO: Basic Linear Algebra Subroutines For Embedded Optimization. *arXiv:1704.02457* 2017.
45. Houska B, Ferreau HJ, Diehl M. An Auto-Generated Real-Time Iteration Algorithm for Nonlinear MPC in the Microsecond Range. *Automatica* 2011; 47(10) : 2279–2285.
46. Quirynen R, Vukov M, Zanon M, Diehl M. Autogenerating Microsecond Solvers for Nonlinear MPC: a Tutorial Using ACADO Integrators. *Optimal Control Applications and Methods* 2014; 36 : 685–704.
47. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends in Machine Learning* 2011; 3(1) : 1–122.
48. Berntorp K. *Particle Filtering and Optimal Control for Vehicles and Robots*. PhD thesis. Department of Automatic Control, Lund University, 2014.
49. Quirynen R, Berntorp K, Di Cairano S. Embedded Optimization Algorithms for Steering in Autonomous Vehicles Based on Nonlinear Model Predictive Control. In: *Proc. American Control Conf. (ACC)*; 2018: 3251–3256.
50. Berntorp K, Quirynen R, Di Cairano S. Steering of Autonomous Vehicles Based on Friction-Adaptive Nonlinear Model-Predictive Control. In: *Proc. American Control Conf. (ACC)*; 2019: 965–970.
51. Berntorp K, Hoang T, Quirynen R, Di Cairano S. Control Architecture Design for Autonomous Vehicles. In: *Proc. IEEE Conf. on Control Techn. and Appl. (CCTA)*; 2018: 404–411.
52. Berntorp K, Inani P, Quirynen R, Di Cairano S. Motion Planning of Autonomous Road Vehicles by Particle Filtering: Implementation and Validation. In: *Proc. American Control Conf. (ACC)*; 2019: 1382–1387.
53. Berntorp K, Quirynen R, Uno T, Cairano SD. Trajectory tracking for autonomous vehicles on varying road surfaces by friction-adaptive nonlinear model predictive control. *Vehicle System Dynamics* 2019; 0(0) : 1–21.
54. Uno T, Quirynen R, Cairano SD, Berntorp K. Vehicle Integrated Control by Nonlinear Model Predictive Control in GNSS Based Autonomous Driving System. *Transactions of Society of Automotive Engineers of Japan* 2019; 50(4) : 1176–1181.
55. Carvalho A, Lefevre S, Schildbach G, Kong J, Borrelli F. Automated driving: The role of forecasts and uncertainty - A control perspective. *European Journal of Control* 2015; 24 : 14–32.

How to cite this article: Quirynen R., and Di Cairano S. (2019), Block-Structured Preconditioning of Iterative Solvers within an Active-Set Strategy for Model Predictive Control Applications, *Optimal Control Applications and Methods*, 2019;00:1–6.