NOVA

University of Newcastle Research Online

nova.newcastle.edu.au

Konnik, Mikhail; De Dona, Jose. "Feasibility of constrained receding horizon control implementation in adaptive optics" IEEE Transactions on Control Systems Technology Vol. 23, Issue 1, p. 274-289 (2015)

Available from: <http://dx.doi.org/10.1109/TCST.2014.2324179>

© 2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Accessed from: <http://hdl.handle.net/1959.13/1056848>

Abstract

Adaptive optics (AO) provides real-time compensation for atmospheric turbulence to improve the resolution of images acquired by ground-based optical telescopes. The actuators in deformable mirrors, which are used as correctors, are constrained by a maximal allowable movement. The control techniques used in current AO systems do not account for these constraints, leading to inferior performance and a risk of damage of the deformable mirror's surface. This article presents a feasibility study for Receding Horizon Control (RHC) with on-line constrained Quadratic Programming (QP). The results of numerical simulations provided in this article are based on realistic models obtained from an optical test-bench. We compare QP algorithms that represent three main methods for convex optimisation: Interior Point, Active Set, and Gradient-based algorithms. It is shown that constrained RHC is computationally feasible for moderate-size adaptive optics systems using hot-started structure-exploiting Active Set QP solvers with bound constraints. An evaluation of performance indicates that RHC is advantageous in terms of atmospheric turbulence rejection in the case of active constraints.

1 Introduction

The ever-moving Earth's atmosphere introduces inhomogeneity in the air refractive index, which leads to blurring of astronomical images acquired by ground-based telescopes. Adaptive optics (AO) can compensate for the atmospheric turbulence in real-time using measurements from wavefront sensors (WFS) to generate control signals for bending the deformable mirror (DM).

Control techniques currently used in AO The most widely used control technique for adaptive optics systems are simple PI [1–3] or PID [4] controllers. Although classical control does not account for the actuators coupling in the deformable mirror and the constraints on maximum actuators movement, PI/PID controllers are still used due to the ease of implementation.

More advanced controllers, such as Linear Quadratic Gaussian (LQG) and minimum variance [5, 6] controllers have been also considered. LQG controllers have been designed [7, 8] and simulated [9, 10] for adaptive optics. The experimental validation of optimal control was shown to outperform classical control [11, 12]. Robust control has been also discussed [13], simulated [14], and tested [15] for adaptive optics applications. Adaptive control for adaptive optics was considered in [16, 17].

The purpose of this study All the aforementioned control techniques are unconstrained, whereas the actuators in deformable mirrors (DM) are obviously constrained by a maximum allowable movement. Too large control signals can cause damage of the mirror's surface. Although there are methods that account for saturation and anti-windup [18, 19], they are not optimisation-based and therefore do not allow to attain the full compensational potential of adaptive optics.

The main motivation for using constrained Receding Horizon Control (RHC) is to provide a control signal that is optimal

within the imposed constraints, which prevents the damage of the surface of a DM and provides better disturbance rejection in constrained operation. Therefore, RHC could make it possible to use adaptive optics even in the case of strong atmospheric turbulence.

The constrained RHC technique remains relatively unexplored in adaptive optics [20–24]. It was used [20] to control a Thermal Deformable Mirror to correct static aberrations in a light beam. The consecutive line searches method [25] in the direction of several Zernike modes was used in [20]. The Nelder-Mead [26] algorithm was used in [21] for correction of optical aberrations in a closed-loop feedback. The Receding Horizon approach was recently applied for control of aero-optical wavefronts [22]

Novelty of this article This article addresses the problem of constrained operation of AO systems in the case of relatively strong atmospheric turbulence (output disturbance), which has not been addressed previously. As a possible solution, constrained Receding Horizon Control with online Quadratic Programming is evaluated. The results of this article indicate that using bound constraints and hot-start for Active Set QP methods makes RHC computationally feasible for control of a moderate-size deformable mirror with computational times within 1 msec with present-day computers.

The simulations are based on models obtained from a hardware optical test-bench [27, 28]. Several variants of coupling between the actuators of the deformable mirror and types of constraints are considered. A comparison of Active Set, Interior Point and Gradient based strategies is provided.

2 Adaptive optics and its Dynamic models

This section briefly describes the components of a typical adaptive optics (AO) system and provides the formulation of dynamic models of AO. An overview of the specific control challenges in adaptive optics, which this article is attempting to address, is presented at the end of this section.

2.1 Brief description of an adaptive optics system

Turbulence of the Earth's atmosphere produces inhomogeneity in the air refractive index that differs for hot and cold air. The volume distribution of the refractive index can be thought of as a surface, or a wavefront. If there is no atmospheric turbulence, the wavefront will be flat, which means that all the light rays have travelled the same distance. Due to mixing of hot and cold air along with strong wind gusts, the wavefront becomes non-flat, which causes the light to travel differently and in turn causes the blurring of astronomical images acquired by ground-based telescopes.

Adaptive optics compensates the atmospheric turbulence in real-time and improves images quality of ground-based telescopes. Typical adaptive optics systems consist of:

1. wavefront sensor (WFS) that measures distortions introduced by atmospheric turbulence;
2. deformable mirror (DM) that corrects instantaneous wavefront distortions;
3. control computer (CC) that tracks a zero reference and generates signals for actuators of the deformable mirror from the wavefront sensor's data.

^{*}Mikhail Konnik is with the School of Electrical Engineering and Computer Science, The University of Newcastle, Callaghan 2308, Australia. mail@mvkonnik.info, mikhail.konnik@uon.edu.au

[†]José De Doná is with the School of Electrical Engineering and Computer Science, The University of Newcastle, Callaghan 2308, Australia. Jose.Dedona@newcastle.edu.au

Figure 1: Components of a typical adaptive optics system.

These components (see Fig. 1) typically operate in closed-loop feedback to compensate for the atmospheric turbulence.

Wavefront measurements are typically obtained from a Shack-Hartmann wavefront sensor (see Fig. 2). Each lens takes a small part of the aperture (sub-aperture) and forms an image of the light source. When the wavefront is distorted, the light spots become displaced from their nominal positions. Displacements of light spots (centroids) are proportional to the average wavefront distortion over the sub-apertures.

The atmospheric turbulence can be represented as an output disturbance and described by a discrete-time state space model. The goal of the controller in the adaptive optics system is to counteract the non-flat (turbulent) wavefront and track a zero reference (that is, a flat wavefront) by bending the continuous reflective surface of the deformable mirror (see Fig. 3) in the direction opposite to the measured wavefront distortion.

2.2 Dynamic models of adaptive optics components

The adaptive optics (AO) system can be effectively modelled by a Linear Time Invariant Multi-Input-Multi-Output system that consists of a wavefront sensor, a controller, and a deformable mirror for the real-time compensation of an output disturbance (atmospheric turbulence). The AO system is subject to time delays (mostly sensor delays) and measurement noise (originating from the photosensor inside the wavefront sensor). The aim of the controller is to track a zero reference, which represents a flat wavefront.

2.2.1 The model of the wavefront sensor

The dynamics of the wavefront sensor (WFS) is described by the integration time T_f of the photosensor, since the output of the sensor is the average of the signal from t to $t+T_f$. An additional delay for data processing and reconstruction is denoted τ_{wfs} . In the discrete-time equivalent, the dynamics of the WFS is usually represented by sensor delays proportional to $T_f + \tau_{wfs}$.

Figure 2: Principle of the Shack-Hartmann wavefront sensor.

Figure 3: The deformable mirror with piezoelectric actuators that control the shape of the surface (uncoated).

2.2.2 The model of the output disturbance

The output disturbance in an adaptive optics system is due to atmospheric turbulence. It was shown [29] that an autoregressive model AR(1) provides limited performance due to the nature of the approximate model. On the other hand, an AR(5) model leads to a loss of performance due to increased sensitivity to variations of temporal characteristics of the atmospheric turbulence. An AR(2) model provides a reasonable agreement with the Kolmogorov model [30] in good seeing conditions and is used in this article in the following form (before time discretisation):

$$G(s)_{i,i}^{\text{atm}} = \frac{1}{s^2 + s \cdot \kappa_1 + \kappa_2}. \quad (1)$$

The poles of the output disturbance model are usually slower than the poles of the DM and in this article we used $\kappa_1 = 2$ and $\kappa_2 = 30$ as mean values for the coefficients of (1).

2.2.3 The model of the deformable mirror

In this article we consider the dynamics of the actuators as a first-order transfer function with a fast pole [31,32], which agrees with experimental data obtained via system identification from the hardware optical test-bench [27]. In order to account for the fact that actuators slightly differ from one to another, the continuous transfer matrix was generated by taking the values of poles from a Normal Distribution $p_{\mathcal{N}} \sim \mathcal{N}(p, \sigma_p^2)$ with mean p and standard deviation $\sigma_p = \alpha \cdot p$. The continuous transfer function for the (i, i) -th actuator is therefore:

$$G(s)_{\text{plant}}^{i,i} = \frac{1}{s + p_{\mathcal{N}}}, \quad (2)$$

where $\alpha = 0.01$ and the poles mean value is $p = 1500$.

Such dynamical model (with coupling between four nearest actuators) was found adequate [27,28] for our target hardware¹, which is the 7×7 Xinetics deformable mirror $\phi 30$ mm in diameter (see Fig. 3) with piezoelectric actuators (37 actuators arranged in a 7×7 square grid, spaced at 5 mm between each actuator).

2.2.4 Coupling between actuators in a deformable mirror

The actuators are arranged in a square grid and attached underneath the reflective surface of a deformable mirror, as seen in Fig. 3. The faceplate of the deformable mirror is made from low-expansion material such as quartz of Corning ULE. When the actuator is energized (pushed up or pulled down), it displaces neighbour actuators due to the sturdiness of the mirror's surface. Therefore, one energized actuator affects neighbour actuators to some extent, making the dynamics of DM more complex.

¹We thank the Research School of Astronomy and Astrophysics at the Australian National University (ANU) for granting access to a Xinetics DM.

Such an influence of an energized actuator upon its neighbours is called *inter-actuator coupling*. There are three cases of coupling between the actuators in a deformable mirror that can be considered, as shown in Fig. 4:

1. no coupling between the actuators, when the energized actuator has no effect on the neighbours (Fig. 4(a));
2. coupling between four nearest neighbour actuators, as seen in Fig. 4(b);
3. coupling between four nearest neighbour and four diagonally adjacent actuators (see Fig. 4(c)).

Figure 4: An illustration of coupling cases for the DM models: (a) decoupled case; (b) coupling between the nearest neighbours; (c) coupling between the nearest neighbour and diagonally adjacent actuators.

The strength of the influence of an actuator on its neighbours depends on properties of the mirror’s surface, inter-actuators spacing and number of actuators: the closer the actuators are to each other, the more the energized actuator affects its neighbours. Actuators coupling can be described as a proportional relationship between the actuators’ transfer functions and described by a coupling degree γ , as explained below. In this case the transfer matrix contains off-diagonal elements $G_{i,j}$:

$$G_{\text{plant}}^{\text{TM}}(s) = \begin{bmatrix} \mathbf{G}_{1,1} & G_{1,2} & 0 & G_{1,4} & 0 & \dots \\ G_{2,1} & \mathbf{G}_{2,2} & G_{2,3} & 0 & G_{2,5} & \dots \\ \vdots & \vdots & \ddots & & & \end{bmatrix}. \quad (3)$$

That is, for this example of 3×3 actuators grid (9 actuators), the coupling TF $G_{5,2}(s)$ between the actuator 5 and actuator 2 is $G_{5,2}(s) = \gamma \cdot G_{5,5}(s)$, where γ is the coupling degree. Continuous-faceplate DMs usually have relatively strong (20–30%, or $\gamma = 0.2$ – 0.3) coupling [33] between the nearest neighbour actuators with inter-actuator spacing 0.2–1 mm and stroke up to $10\mu\text{m}$. In our numerical simulations we consider a coupling degree $\gamma = 0.2$ if not stated otherwise.

2.2.5 Overall model of the system dynamics

The continuous transfer matrices of the plant, $G_{\text{plant}}^{\text{TM}}(s)$, and the disturbance, $G_{\text{atm}}(s)$, were formed using (1) and (2). These transfer matrices were converted to the state space formulation, discretised using a zero-order hold (ZOH) on the inputs with a sample time of $T_s = 10^{-3}$ seconds, obtaining the state space matrices of the overall model. The state space model was augmented to account for both plant and disturbance dynamics:

$$A = \begin{bmatrix} A_{\text{plant}} & 0 \\ 0 & A_{\text{atm}} \end{bmatrix}, B = \begin{bmatrix} B_{\text{plant}} \\ 0 \end{bmatrix}, C = [C_{\text{plant}} C_{\text{atm}}]. \quad (4)$$

This discretised model is used in the following sections for the Receding Horizon Control formulation.

2.3 Motivation and Challenges for constrained Receding Horizon Control in adaptive optics

An adaptive optics system represents a significant challenge from the control point of view. Dynamic matrices for state space formulation are large (ultra-large DMs can have more than 100 ×

100 actuators), and usually augmented with at least two sample delays. The specific challenges are:

1. the control signal has hard constraints that cannot be violated or softened (mirror’s surface might be damaged);
2. specific additional constraints may be required (maximum actuators movement, waffle² mode [34, 35]);
3. limited time for computation of a control input, which must be solved within 1 msec or faster.

This article attempts to address the aforementioned problems using constrained Receding Horizon Control (RHC) with QP algorithms. RHC solves a constrained QP on each sampling instant, producing a control input that is optimal within prescribed constraints. The major problem is that solving QP on-line is computationally demanding. It is worth asking the question “Why do we need optimal control in adaptive optics?”

There were several reasons argued in [8] for **not using** optimal control (LQR or RHC) in adaptive optics:

1. In good seeing conditions a PI control law seems to work well enough.
2. A PI control law is easy to implement.
3. Optimal control usually requires a state estimator and system identification for models of DM and turbulence.
4. The large number of sensors and actuators have made the computational burden infeasible within the millisecond or so duty cycle necessary for adaptive optics.

However, the RHC technique has its own benefits that cannot be underestimated. The main feature of RHC is the ability to generate a control signal that is optimal *within the prescribed constraints*. This allows addressing the following unsolved problems in AO control:

1. respecting hard constraints on the control input prevents surface damage of a deformable mirror;
2. mitigate specific modes, such as a waffle mode [34, 35].

Constrained RHC is considered as a safeguard against potential mirror’s surface damage. Moreover, the ability of RHC to generate control inputs within arbitrary constraints allows utilizing the full potential of a deformable mirror.

3 Mathematical Background of Constrained receding horizon control for Adaptive Optics

While the Quadratic Programming problem in Adaptive Optics is rather simple (small-to-medium scale, convex, relatively well conditioned), the main difficulty is the tight computational deadlines: the constrained QP must be solved within at most 1 msec. A detailed analysis of the control problem in AO leads to the following insights:

1. the coupling between the actuators in a DM affects only close neighbour actuators, therefore the Hessian matrix has a particular *banded* structure (see Subsection 4.2);
2. computational deadlines are very tight: 1 msec at most, therefore one must opt for low-complexity algorithms;
3. the sampling rate in AO is typically very fast (up to 2 kHz), while the output disturbance (atmospheric turbulence) is relatively slow.

²Waffle mode is a movement of two adjacent actuators in opposite directions, which can be hazardous for the mirror’s surface. RHC can be used for waffle mode mitigation via additional constraints in a form of linear inequalities. This is out of scope of this article, but has been considered in [34].

As a result of previous analysis, the following properties of a QP solver are desired:

1. iterations should exploit the structure of the Hessian (see the discussion in Subsection 4.2);
2. hot-start should be very efficient, as the optimisation problem does not change much, in general, from one sampling instance to the next.

Structure exploitation and hot-starting are the two key techniques that allow to achieve extremely fast computation times for AO systems, especially for astronomical applications.

Three different approaches for QP solving are considered in this work: Interior Point, Active Set and Gradient-based. We use the Primal-Dual Interior Point and Dual Active Set algorithms from the QPC toolbox [36], the Parametric Active Set algorithm from qpOASES code [37], CVXGEN [38] implementation of Interior Point, and our own customised implementations of Range-Space Primal Active Set [39], Projected Alternating Barzilai-Borwein [40] and Gradient Projection Conjugate Gradients [41] algorithms. More detailed description of the customised algorithms can be found in [42].

All the algorithms have been implemented in ANSI C and compiled as MEX for MATLAB.

3.1 Receding Horizon Control formulation

Consider the state space model of a linear system:

$$x_{k+1} = Ax_k + Bu_k, \quad y_k = Cx_k + d_k, \quad (5)$$

where x_k is the state, u_k is the control input, y_k is the output, and d_k is a time-varying output disturbance. We use a standard state space RHC formulation [43], where the cost function in vector form can be expressed as:

$$V_{N_p, N_c} = \frac{1}{2} \hat{x}_0^T C^T Q C \hat{x}_0 + \frac{1}{2} X^T Q X + \frac{1}{2} U^T R U, \quad (6)$$

for a state prediction horizon N_p and control prediction horizon N_c , $X = [x_1^T, \dots, x_{N_p}^T]^T$ and $\hat{x}_0$ denotes the *estimate* of the current state obtained with a Kalman filter, and $U = [u_0^T, u_1^T, \dots, u_{N_c-1}^T]^T$ is a vector of future control inputs (here $u_i = 0$ for $i = N_c, \dots, N_p-1$ according to the definition of Γ below). Define the matrices Q and R as follows:

$$Q = \text{diag}\{C^T Q C, \dots, P\} \text{ and } R = \text{diag}\{R, \dots, R\}. \quad (7)$$

where the matrix Q is of size $(nN_p) \times (nN_p)$ and the matrix R is of size $(mN_c) \times (mN_c)$; here m is the number of inputs and n is the number of states. The matrix $Q = I$, the penalty matrix for control energy is chosen as $R = 10^{-12}I$, and the matrix P is chosen as $P = 10^{-2}I$ to match³ the desired performance of RHC. We define matrices Γ and Ω as:

$$\Gamma = \begin{bmatrix} B & \dots & 0 \\ AB & \dots & 0 \\ \vdots & \ddots & \vdots \\ A^{N_p-1}B & \dots & A^{N_p-N_c}B \end{bmatrix}, \quad \Omega = \begin{bmatrix} A \\ A^2 \\ \vdots \\ A^{N_p} \end{bmatrix}.$$

The dynamics in (5) can be expressed over the prediction horizon in vector form as:

$$X = \Gamma U + \Omega \hat{x}_0. \quad (8)$$

Substituting (8) into the cost function (6) we obtain:

$$V_{N_p, N_c} = \bar{V} + \frac{1}{2} U^T \mathbb{H} U + U^T \mathbb{F} \hat{x}_0, \quad (9)$$

where the term $\bar{V}$ is independent of U . The Hessian matrix $\mathbb{H}$ and matrix $\mathbb{F}$ are defined as follows:

$$\mathbb{H} \triangleq \Gamma^T Q \Gamma + R, \quad \mathbb{F} \triangleq \Gamma^T Q \Omega. \quad (10)$$

³We use both P and R matrices as parameters for RHC, and the values $P = 10^{-2}I$ and $R = 10^{-12}I$ provide desired performance in terms of disturbance rejection in our case.

A Quadratic Programming (QP) problem is formulated to minimise the cost function (9) as follows:

$$\begin{aligned} & \underset{U}{\text{minimise}} && \frac{1}{2} U^T \mathbb{H} U + U^T \mathbb{F} \hat{x}_0 \\ & \text{subject to} && M \cdot U \leq \Lambda, \end{aligned} \quad (11)$$

where M is the constraints matrix (we use constraints on the control inputs only), and Λ is the vector of maximum and minimum allowed values for the control input. The Hessian matrix $\mathbb{H}$ is positive definite and hence the constrained solution is unique. The main difficulty, however, is that the QP must be solved within at most 1 msec, which demands a careful comparison of QP strategies and constraints types.

3.2 Gradient Projection Conjugate Gradient

The GPCG algorithm proposed by Moré and Toraldo [41] combines the Gradient Projection (GP) and the Conjugate Gradient (CG) algorithms with switching rules between the two. The GPCG algorithm uses the CG method to explore the face of the feasible region defined by the current iteration, and the GP method to move to a different face. The implementation of GPCG on a parallel architecture is also possible, and results in [44] indicate that GPCG scales well up to 16 processors.

Hot-start in the GPCG algorithm was implemented by feeding the optimal value U_{k-1} from the previous sampling instance as a starting point for the next sampling instance.

3.3 Projected Alternating Barzilai-Borwein

The Barzilai and Borwein (BB) method [40] uses ideas from both steepest descent and Quasi-Newton methods. The Projected Alternating Barzilai-Borwein (PABB) method computes the next step as $U_{k+1} = P_\Psi[U_k - \alpha_k^{\text{BB}} g_k]$, where $\alpha_k^{\text{BB}} > 0$ is a step-length, the vector $g_k = (\mathbb{F} + \mathbb{H}U_k)$ is the gradient at the current value of U_k , and P_Ψ the projection operator onto the feasible set $\Psi = \{U \in \mathbb{R}^m : U_{\min} \leq U \leq U_{\max}\}$ with $U_{\min}$ and $U_{\max}$ denoting the lower and upper bounds of the control inputs. The step-length α_k^{BB} in the PABB algorithm is computed as:

$$\alpha_k^{\text{BB1}} = \frac{s_{k-1}^T s_{k-1}}{s_{k-1}^T y_{k-1}} \quad \text{or} \quad \alpha_k^{\text{BB2}} = \frac{s_{k-1}^T y_{k-1}}{y_{k-1}^T y_{k-1}} \quad (12)$$

where $s_{k-1} = U_k - U_{k-1}$ and $y_{k-1} = g_k - g_{k-1}$. The PABB algorithm has low storage requirements, easy to implement and inexpensive iterations [45]: one iteration of the PABB algorithm has an $O(n)$ complexity.

Hot-start of the PABB algorithm was implemented by feeding the gradient vector g_{k-1} and the optimal value U_{k-1} from the previous sampling instance t_{k-1} as a starting point for the next sampling instance.

3.4 Dual Active Set QP algorithm

The Dual Active Set QP algorithm is based on Goldfarb-Idnani algorithm [46, 47]. Denote $\mathbf{c} = \mathbb{F} \hat{x}_0$ from (11). The algorithm starts with the calculation of the unconstrained minimum $U_0 = -\mathbb{H}^{-1} \mathbf{c}$ of the cost function in (11). When the i -th constraint in $MU \leq \Lambda$ is violated, the algorithm computes a new vector $U_1 = U_0 + p_0$, which satisfies the i -th equality constraint $M_i U_1 = \Lambda_i$, and p_0 is found by solving an equality constrained QP problem:

$$\begin{aligned} p_0 = \underset{p}{\text{argmin}} && (U_0 + p)^T \mathbb{H} (U_0 + p) + \mathbf{c}^T \cdot U_0 + p \\ && \text{subject to} \quad M_i (U_0 + p) = \Lambda_i \end{aligned} \quad (13)$$

The algorithm finds the maximum step length such that the Lagrange multipliers remain strictly positive, $\lambda > 0$.

The implementation of Goldfarb-Idnani algorithm uses fast updates of a QR factorisation [48, 49] via Givens rotations [50] when the constraints are added and removed, which accelerates the algorithm considerably.

3.5 Range-Space Primal Active Set QP

The Range-Space Primal Active Set (RSPAS) algorithm [39] is briefly outlined. Consider $M \in \mathbb{R}^{m \times n}$ is the constraint matrix, and $\Lambda \in \mathbb{R}^{m \times 1}$ is the vector of bounds. At the k -th iteration we are at the feasible iterate U_k and the subset of constraints that are satisfied is represented by $\hat{M}_k \in \mathbb{R}^{z \times n}$, where $0 \leq z \leq n$ is the number of active constraints. The next iteration is computed as $U_{k+1} = U_k + \alpha_k p_k$, where $\alpha_k \in (0, 1]$ is a step-length, and the step direction p_k is found by solving the convex equality constrained sub-problem:

$$\begin{aligned} \min_{p_k} \quad & \frac{1}{2} p_k^T \mathbb{H} p_k + g_k^T p_k \\ \text{subject to} \quad & \hat{M}_k p_k = 0 \end{aligned} \quad (14)$$

where $g_k = (\mathbb{F} + \mathbb{H}U_k)$ is the gradient. The algorithm uses the factorizations $\mathbb{H} = S^T S$ and $\hat{M}S^{-1} = LY^T$, where $S \in \mathbb{R}^{n \times n}$ is the Cholesky factor for $\mathbb{H}$, $L \in \mathbb{R}^{z \times z}$ is lower triangular and $Y \in \mathbb{R}^{n \times z}$ is an orthogonal matrix. Updates of the LQ factorisation of $\hat{M}S^{-1}$ for p_k and λ_k are efficiently computed using classical methods [48, 50].

Hot-start for the RSPAS method uses the previous solution U_{k-1} as an initial starting point for the computation of U_k , and re-uses matrices L_{k-1} and Y_{k-1} from the LQ factorisation at the previous sampling instance t_{k-1} .

3.6 qpOASES Parametric Active Set QP

The Parametric Active Set method, which is used in the qpOASES code [37], builds on the assumption that the active set does not change much from one QP to the next. The general form of the parametric QP can be written as:

$$\begin{aligned} \min_{U} \quad & QP(w_0) = \frac{1}{2} U^T \mathbb{H} U + U^T g(w_0) \\ \text{subject to} \quad & Gx \geq b(w_0). \end{aligned} \quad (15)$$

with a positive definite Hessian matrix $\mathbb{H}$ and a constraint matrix G . To transition from a solved quadratic program $QP(w_0)$ to the next one $QP(w_0^{ew})$, the algorithm moves on a straight line in the parameter space [51], i.e., in the set $\mathbb{P} = \{w_0 \in \mathbb{R}^{n_{w_0}} | QP(w_0) \text{ is feasible}\}$. The qpOASES_sequenceSB routine was compiled to a MEX function and used for hot-started (sequence) control subject to simple bound (SB) constraints and dense matrices. For further details, see [37, 51, 52].

3.7 Mehrotra Primal-Dual Interior Point QP algorithm

In this work we use the QPC's [36] implementation of Mehrotra Predictor-Corrector Primal-Dual QP algorithm [53]. At each iteration, Mehrotra's predictor-corrector algorithm uses Cholesky decomposition to find two different directions for a predictor and a corrector. The Mehrotra's algorithm implementation uses Gondzio's higher order terms [54]. The main idea of the algorithm is to compute an optimizing search direction based on a first order term (this is the predictor). The step size that can be taken in this direction is used to evaluate how much correction to the central path is needed.

3.8 CVXGEN Interior Point QP

To make the comparison of Interior Point methods more fair, we also use the CVXGEN code [38]. CVXGEN uses a standard primal-dual interior point method [55] with Mehrotra predictor-corrector [53]. Each iteration of the primal-dual algorithm requires solution of the KKT matrix and uses a permuted LDL^T factorization [38]. CVXGEN does not use hot start because it does not yield significant improvements in performance.

4 Numerical Analysis of Receding Horizon Control problem

This section provides a brief discussion of the condition number and the structure of the Hessian matrix. The condition number is very important for the gradient methods; however, the structure of the Hessian matrix also affects the performance of Active Set and Interior Point methods. The choice of the prediction horizons for RHC in AO is also discussed.

4.1 Condition number of the Hessian matrix in AO problems

Receding Horizon Control (RHC) in the condensed form uses the two composite matrices from (10), namely $\mathbb{H} \triangleq \Gamma^T \mathcal{Q} \Gamma + \mathcal{R}$ and $\mathbb{F} \triangleq \Gamma^T \mathcal{Q} \Omega$, to solve the QP problem. Therefore, the numerical stability of the solution for the optimal control will depend on the properties of these matrices.

The coupling between actuators in a DM and the variability of actuators' dynamics has an impact on the condition number of the Hessian matrix. This is important because some types of QP solvers (e.g., gradient-based), are sensitive to the conditioning: the more ill-conditioned the Hessian matrix is, the slower is their convergence. The dependency of the condition number on the coupling degree γ between the actuators in the DM is provided in Fig. 5. We used the MATLAB's `cond` command to calculate the l_2 norm condition number κ of the Hessian matrix $\mathbb{H}$ with the shortest horizon $N_p = 2, N_c = 1$.

If the deformable mirror is ideal with perfectly identical actuators with the same dynamics, the only contributor to the condition number is the coupling degree (see Fig. 5). Even for a large coupling degree $\gamma = 0.35$, the Hessian matrix is relatively well-conditioned ($\kappa(\mathbb{H}) < 10^3$).

Figure 5: Condition number of the Hessian matrix versus the coupling degree γ for a deformable mirror with 7×7 actuators for two cases: when the dynamics of actuators is the same, and when the dynamics varies between actuators.

In a real system, however, the actuators are slightly different in their mechanical and electrical properties from one to the next, and therefore the poles of the transfer function of the actuators have different values. One can see that the condition number $\kappa(\mathbb{H}) \sim 10^3$ for the case of similar actuator dynamics, while $\kappa(\mathbb{H}) \sim 10^5$ when the actuators have a variance of $\sigma^2 = 1\%$ in the poles values for the case of coupling degree $\gamma = 0.3$.

Because gradient-based methods use only first-order information, their convergence may deteriorate when the coupling between the actuators (and, consequently, the condition number) is high. We will see this in Subsection 6.3 and 6.4 for the PABB and GPCG algorithms.

4.2 Structure of the Hessian matrix for adaptive optics

In the decoupled case, when one energized actuator does not affect its neighbours (see Fig. 4(a)), the Hessian matrix is diagonal and the QP can be solved within minimal time. The coupling between the nearest neighbour actuators (see Fig. 4(b)) makes the dynamics of the system more complex, which is reflected in the structure of the Hessian matrix (see Fig. 6a). As one can see from Table 1 and also from Fig. 6, the Hessian matrix is sparse (about 80%), positive definite, and has a banded structure that can be exploited by QP solvers.

Such a structure of the Hessian matrix is due to inter-actuators coupling, that is, when one energized actuator affects its neighbour actuators to some extent (see Subsection 2.2.4). Other types of deformable mirrors may have different dynamics and therefore more or less dense Hessian matrices.

Figure 6: The structure of the Hessian matrix for the case of a 7×7 deformable mirror with coupling between nearest neighbour actuators and: a) shortest horizon $N_p = 2, N_c = 1$; b) longer horizon $N_p = 4, N_c = 4$.

The structure of the Hessian matrix can be exploited to speed-up the computations by extracting elements of the Hessian matrix $\mathbb{H}$ to a more compact *banded* matrix $\mathbb{H}_b$. A symmetric $n \times n$ band matrix with K_l subdiagonals and K_u superdiagonals ($K_l = K_u$ for a symmetric matrix) can be stored compactly in a two-dimensional array with $K_u + 1$ rows and n columns. The leading $(K_u + 1) \times n$ part of the banded Hessian matrix $\mathbb{H}_b$ contains the upper triangular part of the symmetric Hessian matrix $\mathbb{H}$, supplied column-by-column [56]. Therefore, computations can be accelerated using standard routines:

- Gradient vector evaluation can be accelerated using the `dsbmv_` BLAS routine
- Objective function evaluation can be accelerated using the `dsbmv_` routine in BLAS.
- Backward and forward substitution can be performed via the `dsbmv_` and `dtbsv_` routines in the BLAS library.

Furthermore, the analysis of the Hessian matrix suggests that *some* small off-diagonal matrix elements can be eliminated by a threshold without noticeable effect on the performance (see [24] for a discussion), although further analysis is outside the scope of this article.

Table 1: Summary of dimensions of optimisation problems in constrained Receding Horizon Control for AO.

Actuators grid $n \times n$	5x5	6x6	7x7	8x8	9x9	10x10
Control inputs	25	36	49	64	81	100
Constraints	25	36	49	64	81	100
Number of states $\hat{x}_0$	75	108	147	192	243	300
Hessian $\mathbb{H}$ sparsity, %	63%	73%	79%	83%	87%	89%
F-matrix sparsity, %	83%	88%	91%	93%	94%	95%

4.3 On the choice of the prediction horizons in RHC for AO

It is well known that long prediction horizons are generally advantageous, assuming sufficient accuracy of the plant model. However, computational deadlines in astronomical adaptive optics are typically very tight, while the plants are very large (plants with 50–100 states and coupling are considered *small*). We consider adaptive optics systems that use continuous-facesheet deformable mirrors with fast piezoelectric actuators and the case of relatively slow atmospheric turbulence.

We therefore choose the shortest prediction horizon $N_p = 2, N_c = 1$ for this work. A comprehensive analysis of different prediction horizons for different algorithms will be the subject of future work (see a brief discussion in Subsection 7.3).

5 Preliminary assessment of Receding Horizon Control applicability for adaptive optics using standard optimisation algorithms

The numerical simulations in this section are intended to provide an initial feasibility analysis of two different QP approaches, namely Active Set [36] and Interior point (Predictor-Corrector [36] and CVXGEN [38]). These algorithms do not use hot-start: we use unmodified QPC and CVXGEN codes. For each approach we consider two cases: general linear inequality constraints and box constraints. The difference is in the implementation: by “general linear inequality constraints” we mean that the matrix $M = \begin{bmatrix} I \\ -I \end{bmatrix}$ and the vector $\Lambda = [U_{max}^T; -U_{min}^T]^T$ are used in (11). By “bound constraints” we mean that the inequalities $U_{min} \leq U \leq U_{max}$ are hardcoded (i.e., no matrix operations are performed). Simulations were based on the system models from Subsection 2.2.

All QP algorithms were compiled as MEX for MATLAB 2007b⁴ using GCC compiler⁵ in Linux. All results are reported for the sampling rate of 1000 Hz if not stated otherwise.

We consider the comparison of computational times for different mirror models with the number of actuators ranging from 5×5 to 10×10 (see Subsection 2.2.3). The coupling degree is assumed to be $\gamma = 0.20$ if not stated otherwise. The percentage of the active constraints is related to the strength of the atmospheric turbulence: the unconstrained case is typical for good seeing conditions, the case of 5% of active constraints may be related to moderate turbulence, and the case of 20% of active constraints corresponds to a strong turbulence.

5.1 Details of Numerical Simulations

All the transfer matrices are converted to state-space and discretised with the sampling rate of 1 kHz, unless stated otherwise. The initial states of both plant and disturbance were assigned to zero. The sequence of the output disturbance was then generated⁶ and re-used for all QP algorithms.

During the simulations, the amplitude of the output disturbance was gradually increasing to activate more and more constraints, and therefore we were able to gather enough data to plot the computational time as a function of the percentage of active constraints. As a consequence of the randomness of the disturbance, the computational time has a distribution that is not known beforehand. Therefore, a scatter plot was made to show the distribution of computational time: *the bigger the markers on the plot, the more probable that the calculations of*

⁴Simulations were run on a Lenovo T420 notebook: Intel Core i7-2640M Processor 2.80GHz with 8 Gb DDR3 SDRAM, Debian GNU/Linux i386 v 6.0 and MATLAB 2007b for UNIX.

⁵The GCC compiler version 4.4.5 was used for the compilation of MEX files: `mex qp.c -lblas -latlas -lm with CFLAGS='-std=c99 -O3 -fPIC -m32'`

⁶We use the model of the output disturbance from (1) driven by a white noise to generate 10000 instances of the disturbance, each instance is a vector of size $(n \times n) \times 1$, where n is the number of actuators.

the control inputs would take a certain time. The computational time was measured by the built-in MATLAB functions `tic-toc`, which have $1\mu\text{s}$ resolution⁷ according to [57].

The spread of computational time can be attributed to many factors. Primarily, the computational time is affected by the randomness of the disturbance process. Also, since the simulations are run on a Linux notebook rather than on an RTOS, the spread of time can be attributed to background processes taking up “wall clock” time. To minimise the impact of background processes, MATLAB was assigned to a designated core of the four-core processor. This puts all QP solvers in similar conditions for the computational time estimation.

5.2 Computational feasibility of RHC: Primal-Dual Predictor-Corrector Interior Point QP algorithm

In this subsection we provide the results of simulations for the QPC’s Interior Point (IP) and the CVXGEN QP code. We compare the feasibility of IP QP for the case of general linear inequality constraints and bound constraints.

5.2.1 Feasibility for general linear inequality constraints case

The model of the deformable mirror with four nearest coupled actuators was used. One can notice from Fig. 7 that the computational time has a very favourable dependence on the percentage of constraints reached: the time grows slowly with the percentage of active constraints. However, the results indicate that the requirement of 1 msec time frame can be satisfied only for a relatively small model with 6×6 nearest coupled actuators. Therefore, the computational performance of the IP QP algorithm with general linear inequality constraints is unacceptable for adaptive optics and we next explore the use of bound constraints for the QP problem.

Figure 7: Comparison of computational time for the Primal-Dual Predictor-Corrector Interior Point QP algorithm with general linear inequality constraints (filled markers on the plot) and bound constraints (empty markers) versus percentage of active constraints (coupling degree $\gamma = 0.20$).

5.2.2 Feasibility for bound constraints case

Expectedly, the performance of the IP QP algorithm with box constraints is considerably better compared to the case of general linear inequality constraints, as seen in Fig. 7. The computational time for the target system of a 7×7 actuators DM

⁷According to the MATLAB documentation [57], `tic-toc` uses `gettimeofday` function in Linux. The output from `gettimeofday` is saved as a `tic` value. When `toc` is called, the `gettimeofday` is called again, the difference is formed with the last `tic` value, divided by a frequency (10^6 Hz) and returned as the elapsed time with a $1\mu\text{s}$ resolution.

decreased from about 1.5 msec to 0.6 msec for the constrained case. The largest size of the problem that can be solved within 1 msec is an 8×8 actuators DM, and in the case of bound constraints the computational time decreased from about 2.5 msec (general linear inequality constraints) to 0.8 msec.

Despite the fact that the speed-up in the case of bound-constrained QP is considerable, there is a noticeable spread⁸ of computational time. The spread is negligible for small problems like 6×6 actuators, but more pronounced for larger problems and can be up to $200\mu\text{sec}$ for an 8×8 actuators DM and 30% of active constraints.

5.2.3 Using faster CVXGEN Interior Point QP code

We also provide the results for the CVXGEN IP QP for the case of bound constraints. Because CVXGEN produces solvers for a single, fixed problem size only, we generated the solver for each size of the DM (from 5×5 to 8×8 actuators) separately. To increase the speed of the solution, we used the duality gap tolerance `settings.eps` = 10^{-2} and the tolerance of norm of the equality and inequality residuals `settings.resid_tol` = 10^{-2} . The results in Fig. 8 show that the CVXGEN code is faster than QPC’s IP algorithm for small problems: it takes only $100\mu\text{s}$ for CVXGEN to solve a 5×5 actuators problem regardless of the active constraints. However, as the optimisation problem grows larger, the benefits of the explicit code generation are reduced.

Figure 8: Computational time for CVXGEN QP IP code versus the percentage of active constraints for different number of actuators in a DM with coupling $\gamma = 0.20$ between nearest neighbour actuators.

Starting from the 7×7 actuators problem, which takes $470 - 560\mu\text{s}$ for the unconstrained case and $650 - 750\mu\text{s}$ for the constrained case, the CVXGEN and QPC’s IP provide similar performance. Larger problems are not suitable for the CVXGEN, as it is designed for small and medium sized problems with up to hundreds of variables [38]. Also, the performance of the CVXGEN does not improve significantly with hot-start [38], which is very important in Adaptive Optics control problems as it will be shown later.

5.2.4 Summary of results for Primal-Dual Predictor-Corrector Interior Point QP

we can conclude that the bound-constrained IP QP algorithm is feasible to control a moderate-size DM with up to 8×8 actuators within 1 msec. An attractive property of the IP QP algorithm is that the computational time is almost the same regardless of the percentage of activated constraints, which may be beneficial for the case of strong turbulence.

⁸We note that the size of the markers in Fig. 7 are related to the probability of the time required to compute the QP solution: the bigger the marker, the more probable that QP solution will take such a time.

5.3 Computational feasibility of RHC: Dual Active Set QP

We provide the comparison results of the general linear inequality and bound constraints for the Dual Active Set (AS) QP algorithm described in Subsection 3.4

5.3.1 Feasibility for general linear inequality constraints

One can see from Fig. 9 that it is possible to solve the constrained QP problem for the deformable mirror with up to 8×8 actuators within 1 msec even for a strong atmospheric turbulence. The computational time of QP solving grows relatively fast with the number of constraints activated. For the example of a 7×7 actuators DM, the time grows from $260\mu\text{sec}$ for the unconstrained case to $720\mu\text{sec}$ for the case of strong turbulence, when 45% of the constraints are active.

The Dual Active Set QP has negligible ($10\text{--}40\mu\text{s}$) spread in computational time even for relatively large problems, as seen in Fig. 9. When only few constraints are reached, even relatively large problems like 9×9 actuators can be computed within a 1 msec time-frame for reasonable (10–15%) percentage of active constraints. Therefore, it is feasible to use the Dual Active Set QP algorithm for Receding Horizon Control for deformable mirrors with up to 9×9 actuators.

Figure 9: Comparison of computational time for the Dual Active Set QP algorithm with general linear inequality constraints (filled markers on the plot) and for the QP algorithm with bound constraints (empty markers) versus percentage of active constraints (coupling degree $\gamma = 0.20$).

5.3.2 Feasibility for bound constraints

While general linear inequality constraints are more flexible, the use of simpler box constraints provides a considerable speed-up: one can see from Fig. 9 that the solution of the constrained QP problem for the target 7×7 actuators deformable mirror is 1.5–2x times faster compared to the general linear inequality constraints case. That is, with bound constraints only one can solve the 7×7 problem within $210\mu\text{sec}$ for the unconstrained case, and about $400\mu\text{sec}$ for the case of strong turbulence (20–30% of active constraints⁹). The maximum size of the problem that can be solved within 1 msec is 10×10 actuators for reasonable (10–15%) percentage of active constraints.

Furthermore, the computational time grows slowly with the percentage of active constraints for the bound-constrained Dual Active Set QP (see Fig. 9) compared to the case of general linear

⁹As before, we want to emphasize that such a case is very pessimistic: the event of constraints violation is rare in real astronomical AO systems. The reason why we use a 20% case is to show that even in such a (highly unlikely) case, the QP can be solved within the specified deadline.

inequality constraints. One can clearly see a considerable speed-up in the case of simplified bound constraints: a 10×10 actuators DM problem can be solved within 1 msec for 10–15% of active constraints. Therefore, even relatively large problems of DM with 10×10 actuators can be solved by Dual Active Set QP within 1 msec for a moderate turbulence.

5.3.3 Summary of results for Dual Active Set QP

The results indicate that Dual Active Set QP is acceptable for RHC in adaptive optics. For example, the 7×7 problem with box-constraints can be solved within $400\mu\text{sec}$ for the case of strong turbulence, and a large problem of 10×10 actuators can be solved within 1 msec for 10–15% of active constraints. Another attractive property is that the computational time grows slowly with the percentage of active constraints for the bound-constrained Dual Active Set QP algorithm.

5.4 Results discussion and further choice of QP algorithms

The results above, while certainly encouraging, lead us to explore other optimisation algorithms. Although Interior Point (IP) methods offer better structure exploiting capabilities, given the typical computational complexity of IP of $\mathcal{O}(n^3)$ and lack of efficient hot-start, they are unlikely to be competitive with simpler gradient-based algorithms with efficient hot-start and typical computational complexity of only $\mathcal{O}(n)$. In AO problems, the condition number is related to the coupling degree that can be significant (see Fig. 5 and the discussion in Subsection 4.1). While gradient-based algorithms suffer from convergence problems in case of poor conditioning of the Hessian matrix, this can be alleviated by preconditioning techniques. Hot-started Active Set methods are very attractive as well, since only few constraints are typically active for a short period of time during the operation of adaptive optics systems. The results of the initial assessment allow to formulate the following focusing areas for subsequent research:

1. active set and gradient-based algorithms are likely to be the main candidates for QP solvers in adaptive optics;
2. we have to resort to bound constraints in order to satisfy the tight computational requirements in AO;
3. hot-start and structure-exploiting must be used (and be efficient) to further accelerate the QP solution.

To meet the tight deadlines of AO control, we need to achieve a trade-off: we sacrifice flexibility (general linear inequality constraints) for additional speed (bound constraints). Therefore, *in the following sections we will consider only bound constraints to meet extremely tight computational time requirements.*

6 Feasibility of Constrained Receding Horizon Control for adaptive optics with customised optimisation algorithms

To exploit the banded structure of the Hessian matrix and accelerate the solution of a QP problem given tight requirements on computational time (less than a millisecond), we implemented three additional algorithms: Range-Space Primal Active Set (QP RSPAS), Projected Alternating¹⁰ Barzilai-Borwein (PABB), and Gradient Projection Conjugate Gradient (GPCG and its preconditioned version GPpreconCG). Additional details on the implementations can be found in [42]. We also used the qpOASES code to compare with QP RSPAS. The results of the simulations are provided in this section.

¹⁰The PABB algorithm represents a very promising family of Alternating Projections algorithms, which is why PABB is chosen by the authors over other algorithms, e.g., Nesterov's Fast Gradient Method [58].

6.1 Results of Numerical simulations for Range-Space Primal Active Set QP algorithm

This subsection provides the results for the Range-Space Primal Active Set (RSPAS) QP algorithm, each iteration of which is feasible (i.e., can be stopped and the intermediate result will be within the feasible set). This algorithm uses a classical procedure for LQ updates [48, 49], can be efficiently hot-started, exploits the problem's structure (banded matrices and simple bound constraints) and its physics (slow disturbance dynamics and fast sampling).

Figure 10: Computational time for Range-Space Primal Active Set QP algorithm (with bound constraints only) versus percentage of active constraints in a hot-start operation mode (coupling $\gamma = 0.20$ between the nearest actuators only).

6.1.1 Feasibility of Range-Space Primal Active set algorithm for bound constraints case

The hot-start allows exploiting the inherently fast sampling rate in adaptive optics systems and provides a considerable speed-up over a cold-start operation. The Primal Active Set QP in cold-start mode solves a 7×7 actuators DM problem within $600\mu\text{s}$ when 25% of the constraints are active, while in hot-start mode it takes only about $40\mu\text{s}$ (see Fig. 10). Even large systems like 10×10 actuators can be solved within $150\mu\text{s}$ in hot-start mode for the pessimistic case of 20–40% of active constraints (see Fig. 10).

Figure 11: Computational time for Range-Space Primal Active Set QP algorithm versus the number of actuators (coupling degree $\gamma = 0.20$) in the model of a deformable mirror: no active constraints ($\diamond$), case when 5% of constraints are active (∇), and case when 20% of constraints are active ($\circ$). Markers have been slightly split apart for illustrative purposes.

6.1.2 Growth of computation time with size of the problem

The comparison of growth rate of computation time is provided in Fig. 11 for the case of 1000 Hz sampling rate. The computational time for the hot-started Active Set algorithm grows slowly and is within the time-frame of $80\mu\text{s}$ even for large (10×10 actuators DM) systems and 20% of active constraints. As before, we note that such an operation is highly unlikely for adaptive optics systems, but it is useful to know the behaviour of the algorithm under these severe conditions.

6.1.3 Summary of results for Range-Space Primal Active set QP

The hot-start operation for Range-Space Primal Active Set is crucially important and advantageous, especially in the case when many constraints are active. Hot-start operation for the same 7×7 problem takes only about $40\mu\text{s}$ almost regardless of the percentage of the constraints. Even large systems like 10×10 actuators can be solved within $150\mu\text{s}$ for the pessimistic case of 20–40% of active constraints in hot-start mode, while in cold-start it takes longer than $1000\mu\text{s}$. The case of few active constraints takes about $70\mu\text{s}$ to solve for Primal Active Set. Hot-start is less beneficial for slower sampling rates such as 250Hz but still advantageous.

6.2 Results of Numerical simulations for qpOASES Active Set

This section provides the results of numerical simulations for the `qpOASES_sequenceSB` code [37] that is hot-started and uses simple bound (SB) constraints.

Figure 12: Computational time versus the percentage of active constraints for `qpOASES_sequenceSB` (coupling $\gamma = 0.20$ between nearest neighbour actuators, hot-start, bound constraints).

6.2.1 Feasibility of qpOASES for bound constraints case

The results for the case of nearest neighbours coupled actuators are provided in Fig. 12. In hot-start mode, qpOASES uses only one iteration to find a new optimum. Furthermore, the computational time does not grow with the percentage of active constraints. The computational time for the target DM of 7×7 actuators is about $55\text{--}60\mu\text{s}$, and stays almost the same for 5% (moderate turbulence) and 20% active constraints (strong turbulence). Large control problems such as 10×10 actuators DM can be solved within $140\text{--}180\mu\text{s}$ even for the pessimistic case of 20–40% of active constraints in the hot-start mode.

6.2.2 Growth of computation time with size of the problem

We can make a comparison between the computational performance of the qpOASES in Fig. 13 and the Range-Space Primal Active Set algorithm in Fig. 11. One can see that the Range-Space algorithm is faster for large problems (8×8 actuators DM) and similar for smaller problems like 6×6 actuators DM.

Our Range-Space Primal Active Set (RSPAS) QP implementation is faster than qpOASES in the unconstrained case, mainly due to the fact that the RSPAS algorithm uses pre-computed matrices (e.g., inverted Hessian). This is reasonable, since for Adaptive Optics systems the Hessian matrix does not change, and the main operational mode is unconstrained. In the constrained case RSPAS is 1.3–2x times faster than qpOASES, which again should not be surprising: qpOASES uses primal and dual variables within the parameterisation, while RSPAS is a Primal algorithm.

Figure 13: Computational time versus the number of actuators (coupling degree $\gamma = 0.20$) in the model of a deformable mirror: qpOASES code with simplified bounds (SB).

6.2.3 Summary of results for qpOASES

In hot-start mode qpOASES uses only one iteration to find a new optimum, therefore it takes only $55\text{--}60\mu\text{s}$ for a 7×7 DM with nearest coupled actuators and $140\text{--}180\mu\text{s}$ for a 10×10 DM. An attractive feature of qpOASES is that the computational time does not grow with the percentage of active constraints. The comparison of qpOASES against RSPAS QP shows that the RSPAS algorithm is 1.3–2x times faster for the case of large problems and few active constraints. This is mainly because the RSPAS algorithm uses pre-computed matrices.

6.3 Results of Simulations for PABB QP algorithm

The Projected Alternating Barzilai-Borwein (PABB) algorithm is fast and easy to implement. However, like other gradient-based algorithms, the convergence of PABB depends on the condition number of the Hessian matrix and the tolerance τ . Hot-start for PABB is advantageous: the acceleration is more pronounced for the unconstrained case (up to 2–3 times less iterations), and slightly less beneficial for the constrained case (up to 1.5–2.5 times less iterations). Therefore, all the results below are for hot-started PABB.

6.3.1 Feasibility of PABB algorithm for bound constraints case

As the coupling between the actuators decreases, which corresponds to a decrease of the condition number of the Hessian,

the convergence and the computational time of the PABB algorithm improves. This is shown in Fig. 14 for a coupling degree $\gamma = 0.20$, which corresponds to a condition number $\kappa(\mathbb{H}) \sim 10^2$ (see Fig. 5). Expectedly, in the case of acceptable tolerance (such as $\tau = 10^{-3}$) and good conditioning the PABB algorithm is very fast, as seen in Fig. 14. Even a large optimisation problem, such as 10×10 actuators DM, takes $95\text{--}110\mu\text{s}$ for a tolerance of $\tau = 10^{-3}$, the case of 5% active constraints (moderate turbulence) takes $70\text{--}90\mu\text{s}$, and the 20% active constraints case takes $80\text{--}90\mu\text{s}$.

6.3.2 Growth of computation time with size of the problem for PABB QP algorithm

To illustrate the dependence of the PABB algorithm's convergence on the inter-actuator coupling, we provide the comparison of two cases: heavily coupled actuators DM $\gamma = 0.25$ (poor-conditioned Hessian, $\kappa(\mathbb{H}) \sim 10^3$) and a case of weaker coupling $\gamma = 0.20$ between nearest actuators (well-conditioned Hessian, $\kappa(\mathbb{H}) \sim 40$). The tolerance is set to $\tau = 10^{-3}$. The results are provided in Fig. 15.

6.3.3 Summary of results for PABB QP algorithm

PABB can be efficiently hot-started by feeding back the previous optimal solution and gradient vector as a starting point. The acceleration is more beneficial for the unconstrained case with up to 2–3x times speed-up. The convergence rate of PABB gets worse as the coupling degree (and, consequently, the condition number of the Hessian matrix) increases. For problems larger than 8×8 actuators, the computational time grows rapidly. Therefore, PABB QP is more suitable for relatively well-conditioned problems, where $\kappa(\mathbb{H}) \leq 10^3$. These circumstances compel us to explore other algorithms in the subsequent subsections, which allow both projection and preconditioning.

6.4 Results of Numerical simulations for Gradient Projection Preconditioned Conjugate Gradient GPCG QP algorithm

The GPCG QP algorithm (and its preconditioned conjugate gradient version GPpreconCG) uses two stages to solve the bound-constrained optimization problem: the Gradient Projection (GP) stage for the constrained part of the problem, and the Conjugate Gradient (CG) stage for the unconstrained part of the problem. We use the incomplete Cholesky factorisation as a preconditioner for the GPpreconCG algorithm.

Figure 14: Computational time for the PABB QP versus the percentage of active constraints for different number of actuators in a DM with coupling $\gamma = 0.20$ between nearest neighbour actuators (tolerance $\tau = 10^{-3}$, hot-start).

There are four main parameters for the GPCG algorithm: for the sufficient progress test (η_2 for the GP part and η_1 for the CG part) and for the sufficient decrease condition (μ_2 for GP and μ_1 for CG). The parameters η and μ are described in detail in [41]. Our simulations suggest the parameters can be chosen as $\eta_1 = 0.25$ and $\mu_1 = 0.25$ for the CG part and $\eta_2 = 0.60$ and $\mu_2 = 0.25$ for the GP part. The tolerance $\tau = 10^{-3}$ is chosen for both GPCG and GPpreconCG algorithm.

6.4.1 Feasibility of GPCG algorithm for bound constraints case

Our simulations show that GPCG can be accelerated by 1.2–1.5x with hot-starting and even more so with preconditioning (we use incomplete Cholesky factorisation). This is why we implemented GPpreconCG, where *preconCG* is the Preconditioned Conjugate Gradient.

The results for the GPpreconCG (see Fig.16) are encouraging. The unconstrained case takes only $50\mu\text{s}$ with GPpreconCG, as opposed to $200\mu\text{s}$ for the GPCG. The case of moderate turbulence, when 5% of the constraints are active, requires about $80\mu\text{s}$ for GPpreconCG — a considerable improvement over GPCG with $170\mu\text{s}$. For larger 10×10 DM, improvements are more pronounced: GPpreconCG requires $95\mu\text{s}$ in the unconstrained case, $130\text{--}150\mu\text{s}$ in the case of 5% active constraints, and $200\text{--}220\mu\text{s}$ in the case of 20% active constraints. This corresponds to 7x speed-up for the unconstrained case and 3x for the case of few active constants.

(a)

(b)

Figure 15: Computational time versus the size of the optimization problem (number of actuators in a DM) for PABB QP with tolerance $\tau = 10^{-3}$ for: (a) coupling degree $\gamma = 0.25$; (b) coupling degree $\gamma = 0.20$.

Figure 16: Computational time versus the percentage of active constraints for different number of actuators in a deformable mirror with coupling $\gamma = 0.20$ between nearest neighbour actuators, hot-started GPpreconCG algorithm.

6.4.2 Growth rate of computation time with size of the problem

A comparison of computational time growth rate for both cases of coupling is presented in Fig. 17 for the GPpreconCG algorithm. It is evident that the computational time for GPpreconCG grows linearly with the size of the problem, which is a highly desirable feature of a QP algorithm for adaptive optics systems. The algorithm is therefore suitable for AO control problems, especially considering the additional attractive feature of good multiprocessor scaling [44].

Figure 17: Growth rate of the computational time for the GPpreconCG QP algorithm for nearest actuators coupling (filled markers) and for coupling between both nearest and diagonal adjacent actuators (empty markers) versus the size of the deformable mirror (coupling degree $\gamma = 0.20$).

6.4.3 Summary of results for GPCG and GPpreconCG

We observed a considerable speed-up in both unconstrained and constrained cases for the GPCG QP when preconditioning was used. The computational time for GPpreconCG grows linearly with the size of the problem, which is a desirable feature of a QP algorithm for adaptive optics systems. The GPpreconCG algorithm, however, is slower than PABB QP, since GPpreconCG evaluates the objective function to monitor the progress in both parts (CG and GP). Nonetheless, GPCG (and especially its preconditioned version, GPpreconCG) can be regarded as suitable for adaptive optics problems.

7 Receding Horizon Control and LQR: a Comparison and Performance evaluation

The results of the previous sections indicate that RHC is feasible for adaptive optics applications. This section provides a brief performance analysis of the RHC approach and a comparison with established methods such as LQR control.

7.1 Output disturbance rejection performance for Receding Horizon Control and LQR control

The same model of a deformable mirror with 15×15 nearest coupled actuators was used for both RHC and infinite horizon LQR, and the output disturbance was generated once and then re-used (the setup for the numerical simulations was the same as described in Subsection 5.1, footnote ⁶). Then both controllers were run to compensate the same output disturbance. The sum of squares of the residual disturbance at each actuator (i.e., the residual output error, or remaining disturbance) was used for the performance comparison. We considered three different cases of LQR, namely high-gain control (when the matrix $R = 10^{-12}I$), low-gain control (the matrix $R = 10^{-7}I$) and an “optimised” (medium-gain) case when $R = 10^{-8}I$.

7.1.1 High and Low gain LQR versus RHC

The matrix $R = 10^{-12}I$ is considered for high-gain control for both infinite horizon LQR and RHC (the shortest horizon $N_p = 2$, $N_c = 1$ was used). Both controllers rejected the same output disturbance, and the distribution of residuals was recorded for further comparison. The plant is a Deformable Mirror (DM) with 225 actuators¹¹ arranged in a square 15×15 grid, and each actuator is coupled with the four nearest neighbour actuators with a coupling degree $\gamma = 0.2$ if not stated otherwise. The comparison between the RHC and clipped LQR with different gains is provided in Fig. 18. Here the control inputs are shown for the same actuator 16 in the 15×15 actuators DM.

Figure 18: Control inputs for constrained RHC and clipped LQR for different gains: low-gain, high-gain and optimised gain cases. The data are for the actuator 16, which is coupled ($\gamma = 0.20$) with its nearest neighbours.

The choice of matrix $R = 10^{-12}I$ yields a high-gain control, and the coupling between the actuators makes the problem of over-saturation apparent. One can see in Fig. 18 (marked by $-\bullet-$) that the high-gain LQR control is problematic for DM operation: abrupt and rapid changes in control inputs are hazardous for the mirror’s surface. The low-gain LQR control ($R = 10^{-7}I$, marked by $-\blacksquare-$), on the other hand, does not

present the problem of being too aggressive, but it does not use the full control power of the actuators. The comparison shows that even for the optimised gain in clipped LQR, there are still some spikes and abrupt changes in the control input (see in Fig. 18, the line marked by $-\blacktriangle-$), unlike for RHC control (marked by $-\odot-$).

7.1.2 Comparison of output disturbance rejection performance for LQR and RHC

Different gains in infinite horizon LQR produce more or less aggressive controls, and, consequently, different levels of residual output disturbance. The comparison in terms of *sum of squares* of residuals¹² in Fig. 19 indicates that RHC in the constrained case performs better in terms of disturbance rejection performance compared to saturated LQR.

Figure 19: Sum of squares of residual disturbance for a deformable mirror with 225 (15×15) actuators: comparison between the unconstrained (from 0 to 180) and constrained (from 180 to 600) operation for LQR and RHC. The case of nearest actuators coupling is presented (coupling degree $\gamma = 0.20$).

It is apparent from Fig. 19 that while the high-gain LQR (marked by $-\bullet-$) provides similar disturbance rejection performance in the unconstrained case as RHC (marked by $-\odot-$), LQR is clearly the worst when the control inputs are saturated¹³. Optimised gain LQR (marked by $-\blacktriangle-$) provides almost the same disturbance rejection as RHC in the constrained case, but it is worse at cancelling atmospheric turbulence in the unconstrained case, as one can see in Fig. 19. The low-gain control (marked by $-\blacksquare-$), on the other hand, does not use the full control power of the actuators and therefore provides a poor disturbance rejection performance. Therefore, simply clipping (saturating) the control inputs in LQR leads to poor output disturbance performance, while RHC (marked by $-\odot-$) provides the best choice of control inputs.

7.2 Comparison of Receding Horizon Control and LQR control in terms of computational time

It is expected that RHC would be significantly slower than LQR — after all, to compute the control inputs via LQR one has to multiply two matrices (or a matrix by a vector). However, if hot-starting, structure exploiting and simple bound constraints are used in the QP algorithm for the Adaptive Optics control problem, the difference between LQR and RHC is not as dramatic as one would have expected.

To evaluate the computational performance of LQR and compare it to RHC, a simple ANSI C / MEX routine was written

¹¹We use a larger model of the DM to get more accurate results for the output disturbance rejection performance comparison between LQR and RHC in terms of *sum of squares* of residuals (remaining uncompensated disturbance).

¹²We refer to the sum of squares Θ of (uncompensated) residuals for all actuators in a deformable mirror at a given time slice, i.e., $\Theta = \sum_{i=1}^{225} (\sigma_i^2)$.

¹³The output disturbance in these simulations starts from zero, and the amplitude is being gradually increased to reach more and more constraints as time progresses, in the same way as described in Subsection 5.1.

that computes $U_{new} = -K_{gain}\hat{x}_0$ and projects (by clipping the control inputs) the values onto the box-constrained feasible set $\Psi = \{U \in \mathbb{R}^m : U_{min} \leq U \leq U_{max}\}$, with U_{min} and U_{max} denoting the lower and upper bounds of the control inputs. The computational time performance of LQR for constrained and unconstrained cases was measured similarly to the QP algorithms in Sections 5-6. Expectedly, the computational time of LQR is faster: it takes about 25–30 μ s to compute the control inputs for a 7×7 actuators DM, and 60–70 μ s for a 10×10 actuators DM, providing a control rate of about 16 kHz even for a large mirror with 100 actuators.

Figure 20: Comparison between LQR and hot-started RHC with Range Space Active Set QP in AO: dependency of the computational time versus percentage of active constraints. Cases of 5×5 (25 actuators), 7×7 (49 actuators) and 10×10 (100 actuators) are shown for a coupling degree $\gamma = 0.20$.

However, the dynamics of the atmospheric turbulence is usually considerably slower compared to the dynamics of the actuators in a deformable mirror, and the frame-rate of a wavefront sensor usually is 500–1000 Hz or more. This fact can be conveniently exploited by optimisation methods via hot-start to solve the QP problem online: a QP algorithm can quickly (within few iterations) locate the new optimum even in cases of many active constraints.

Figure 21: Computational time for LQR and RHC with Range Space Active Set QP versus the number of actuators (coupling degree $\gamma = 0.20$) in a deformable mirror.

Therefore, it is not surprising that *under the aforementioned conditions* the Receding Horizon Controller with hot-started Range-Space Primal Active Set QP (see Subsection 6.1) is almost as fast as the LQR controller, as seen in Fig. 20. The QP solution, however, is slightly slower than LQR, especially when more constraints are active, which corresponds to a strong tur-

bulence (15–20% and more). For example, LQR takes 60–65 μ s for a 10×10 actuators DM, while RHC takes from 65–70 μ s to 90–120 μ s. This is because two or more iterations are necessary to find the new constrained optimum, which can also be seen in Fig. 21 that shows the growth rate of computational time for RHC and LQR.

7.3 The length of the prediction horizon and the computational time for Receding Horizon Control

The control input in astronomical adaptive optics must be computed as fast as possible: the deadline is typically 1 msec or faster. This poses the question: “What length of the prediction horizon is computationally feasible for adaptive optics?”

This feasibility study covers a typical case of an adaptive optics system, which uses a continuous-facesheet deformable mirror with fast actuators (actuators’ dynamics are assumed to be far faster than the dynamics of the turbulence).

Table 2: Influence of prediction horizon length in constrained Receding Horizon Control for AO on the computational time and residual disturbance (7×7 actuators DM, RSPAS QP).

Horizon N_p, N_c	Hessian size $n \times n$	Computational time active constraints, %		Sum of squares of residuals for	
		0% active	5% active	0% active	5% active
2, 1	49x49	28 μ s	33 μ s	$1.24 \cdot 10^{-6}$	$2.53 \cdot 10^{-3}$
2, 2	98x98	30 μ s	345 μ s	$0.87 \cdot 10^{-6}$	$2.50 \cdot 10^{-3}$
3, 3	147x147	35 μ s	2043 μ s	$0.89 \cdot 10^{-6}$	$2.58 \cdot 10^{-3}$
4, 4	196x196	44 μ s	4921 μ s	$0.87 \cdot 10^{-6}$	$2.52 \cdot 10^{-3}$
5, 5	245x245	47 μ s	5287 μ s	$0.88 \cdot 10^{-6}$	$3.10 \cdot 10^{-3}$

The numerical simulations presented in Table 2 indicate that for the *considered* case, longer horizons do not provide noticeable benefits: while the disturbance rejection slightly improved from $1.24 \cdot 10^{-6}$ to about $0.9 \cdot 10^{-6}$, the computational time for the constrained case and RSPAS QP algorithm grows quickly with the increase of the prediction horizons.

One should note, however, that other types of deformable mirrors may have different dynamics and therefore require longer prediction horizons. For example, for the AO system with a membrane mirror considered in [22, 23], the prediction horizon of $N_c = 20$ was used. Further analysis of different prediction horizons is out of scope of this article, but will be the subject of future work.

8 Discussion and Conclusion

Adaptive optics provides real-time compensation for atmospheric turbulence using measurements from a wavefront sensor to generate control signals for bending the deformable mirror. The actuators in the deformable mirror are constrained by a maximal allowable stroke. The control techniques currently used in adaptive optics do not account for these constraints. The study in this article has evaluated the feasibility of using Receding Horizon Control to handle these constraints.

8.1 Computational feasibility of Receding Horizon Control

Gradient-based methods are attractive because of ease of implementation, effective hot-start, lower computational complexity of $\mathcal{O}(n)$, and the possibility of preconditioning. The PABB QP algorithm is suitable only for relatively well-conditioned problems, where $\kappa(\mathbb{H}) \leq 10^3$. The GPpreconCG algorithm, which allows both projection and preconditioning, seems more promising: it allows solving the target 7×7 DM problem within about 80 μ s for few active constraints, and for larger problems of 10×10 DM it requires about 100 μ s.

While Interior Point QP algorithms are unlikely to be competitive due to a lack of an efficient hot-start and a typical complexity of $\mathcal{O}(n^3)$, they might be useful in the case of strong atmospheric turbulence and longer prediction horizons since their run-time is relatively constant in comparison to the number of constraints active at a given optimal solution.

The most promising algorithms for RHC in adaptive optics appear to be the Primal (RSPAS) and the Parametric (qpOASES) Active Set. Using fast factorisation updates, bound constraints and hot-start, we can considerably reduce the computational time and complexity (each iteration is $\mathcal{O}(n^2)$ flops). For the RSPAS algorithm, it takes only about $40\mu\text{s}$ to solve a hot-started QP for a deformable mirror with 7×7 coupled actuators and 1000 Hz sampling rate, regardless of the percentage of active constraints. The largest problem of a 10×10 actuators DM, can be solved within $150\mu\text{s}$ even for a pessimistic case of 20-40% of active constraints in hot-start mode. This is well within the requirement of 1 msec deadline and is almost as fast as a simple Linear Quadratic Regulator. The results of the simulations show that *RHC clearly outperforms clipped LQR* in terms of atmospheric turbulence rejection performance at the price of a slightly increased computational time.

We conclude that in cases of fast sampling rates and relatively slow turbulence, which is common for astronomical adaptive optics, one can use Receding Horizon Control with online QP for both unconstrained and constrained cases *with similar computational time compared to LQR*.

8.2 Distributing the computational load

From a practical perspective, problems larger than 10×10 actuators should be sent to different CPU cores. Given current high-performance computers, such as the IBM System x3750 with 4x Intel Xeon Processors that have 8 cores 2.4GHz each, we can handle $4 \times 8 \times 100 = 3200$ actuators, allowing us to deal with mirrors of up to 56×56 actuators. In an effort to be able to use Receding Horizon Control on large telescopes, where the deformable mirrors have 60×60 actuators and more, we plan to investigate in future work the parallelisation of the required optimisations via Distributed Receding Horizon Control.

8.3 Conclusion

Thus, our conclusion is that constrained Receding Horizon control is computationally feasible to be used in adaptive optics systems. Despite very tight computational time requirements, we can exploit the banded structure of the Hessian matrix, hot-start and bound constraints to make RHC feasible for adaptive optics control problems. This allows the use of adaptive optics in cases of stronger atmospheric turbulence and therefore constitutes an important step towards attaining the full compensation potential of deformable mirrors.

8.4 Acknowledgements

The authors would like to thank the three anonymous reviewers for their thoughtful and insightful comments that helped to considerably improve the clarity of this manuscript.

References

- [1] C. Dessenne, P. Madec, and G. Rousset, “Modal prediction for closed-loop adaptive optics,” *Optics letters*, vol. 22, no. 20, pp. 1535–1537, 1997.
- [2] E. Gendron and P. Léna, “Astronomical adaptive optics. II. Experimental results of an optimized modal control.” *Astronomy and Astrophysics Supplement Series*, vol. 111, p. 153, 1995.
- [3] C. Dessenne, P. Madec, and G. Rousset, “Optimization of a predictive controller for closed-loop adaptive optics,” *Applied optics*, vol. 37, no. 21, pp. 4623–4633, 1998.
- [4] A. Enmark, T. Berkefeld, O. von der Lühe, and T. Andersen, “Simulation of adaptive optics for the vacuum tower telescope,” *Experimental Astronomy*, vol. 21, no. 2, pp. 87–99, Apr. 2006.
- [5] D. Looze, “Minimum variance control structure for adaptive optics systems,” *JOSA A*, vol. 23, no. 3, pp. 603–612, 2006.
- [6] C. Kulcsár, H. Raynaud, C. Petit, and J. Conan, “Minimum variance prediction and control for adaptive optics,” *Automatica*, vol. 48, no. 9, pp. 1939 – 1954, 2012.
- [7] D. Wiberg, C. Max, and D. Gavel, “A spatial non-dynamic LQG controller: Part II, theory,” in *43rd IEEE Conference on Decision and Control, Nassau, Bahamas, 14-17 December*, vol. 3, 2004, pp. 3333–3338.
- [8] —, “A spatial non-dynamic LQG controller: Part I, Application to adaptive optics,” in *43rd IEEE Conference on Decision and Control, Nassau, Bahamas, 14-17 December*, vol. 3, 2004, pp. 3326–3332.
- [9] D. Looze, M. Kasper, S. Hippler, O. Beker, and R. Weiss, “Optimal compensation and implementation for adaptive optics systems,” *Experimental Astronomy*, vol. 15, no. 2, pp. 67–88, 2003.
- [10] B. Le Roux, J. Conan, C. Kulcsár, H. Raynaud, L. Mugnier, and T. Fusco, “Optimal control law for classical and multiconjugate adaptive optics,” *Journal of the Optical Society of America A*, vol. 21, no. 7, pp. 1261–1276, 2004.
- [11] C. Petit, J. Conan, C. Kulcsár, H. Raynaud, T. Fusco, J. Montri, F. Chemla, and D. Rabaud, “Off-axis adaptive optics with optimal control: experimental and numerical validation,” in *Proceedings of SPIE, volume 5903, pages 59030P*, 2005.
- [12] C. Petit, J.-M. Conan, C. Kulcsár, H.-F. Raynaud, T. Fusco, J. Montri, and D. Rabaud, “Optimal control for multi-conjugate adaptive optics,” *Comptes Rendus Physique*, vol. 6, no. 10, pp. 1059 – 1069, 2005.
- [13] N. Doelman, “High performance AO control for time-variant wavefront distortions,” in *Proceedings of SPIE – Astronomical Telescopes and Instrumentation, Adaptive Optics Systems II*, vol. 8447-108, 2012.
- [14] K. Hinnen, M. Verhaegen, and N. Doelman, “A data-driven H₂-optimal control approach for adaptive optics,” *IEEE Transactions on Control Systems Technology*, vol. 16, no. 3, pp. 381–395, 2008.
- [15] —, “Adaptive optics H₂-optimal control design applied on an experimental setup,” in *Proceedings of SPIE – Advances in Adaptive Optics II*, vol. 6272, 2006, pp. 62722S–1.
- [16] Y. Liu and S. Gibson, “Adaptive optics with adaptive filtering and control,” in *Proceedings of the American Control Conference, volume 4, pages 3176–3179*, IEEE, 2005.
- [17] N. Arancibia, N. Chen, S. Gibson, and T. Tsao, “Adaptive control of a MEMS steering mirror for suppression of laser beam jitter,” in *Proceedings of the American Control Conference*. IEEE, 2005, pp. 3586–3591.
- [18] J.-P. Véran and L. A. Poyneer, “Strategies for dealing with DM actuator saturations in advanced AO systems,” *Adaptive Optics: Analysis and Methods/Computational Optical Sensing and Imaging Topical Meeting*, p. AWB3, 2007.

- [19] M. Kasper, E. Fedrigo, D. P. Looze, H. Bonnet, L. Ivanescu, and S. Oberti, "Fast calibration of high-order adaptive optics systems," *J. Opt. Soc. Am. A*, vol. 21, no. 6, pp. 1004–1008, Jun 2004.
- [20] M. Boer, K. Hinnen, M. Verhaegen, R. Fraanje, G. Vdovin, and N. Doelman, "Control of a thermal deformable mirror: Correction of a static disturbance with limited sensor information," in *Adaptive Optics for Industry and Medicine*, vol. 102, 2005, pp. 61–71.
- [21] G. V. Vdovin, "Optimization-based operation of micromachined deformable mirrors," in *Proceedings of SPIE - Adaptive Optical System Technologies*, vol. 3353, no. 1, 1998, pp. 902–909.
- [22] J. Tesch and S. Gibson, "Receding-horizon adaptive control of aero-optical wavefronts," in *Proceedings of SPIE - Acquisition, Tracking, Pointing, and Laser Systems Technologies XXVI*, vol. 8395, 2012, pp. 83 950B–1.
- [23] J. Tesch, S. Gibson, and M. Verhaegen, "Receding-horizon adaptive control of aero-optical wavefronts," *Optical Engineering*, vol. 52, no. 7, p. 071406, 2013.
- [24] Mikhail V. Konnik, José De Doná, and James Stuart Welsh, "On application of constrained receding horizon control in astronomical adaptive optics," in *Proceedings of SPIE Astronomical Telescopes + Instrumentation - Adaptive Optics Systems III*, vol. 8447-110, 2012.
- [25] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.
- [26] J. A. Nelder and R. Mead, "A simplex method for function minimization," *The computer journal*, vol. 7, no. 4, pp. 308–313, 1965.
- [27] A. Smith, "Adaptive optics: Final year project - electrical engineering," the University of Newcastle, Australia, Tech. Rep., 11 December 2010.
- [28] A. Beasley, "Xinetics deformable mirror," Research School of Astronomy and Astrophysics (RSAA), Australian National University, Mount Stromlo, Tech. Rep., 2007.
- [29] C. Petit, S. Meimon, T. Fusco, C. Kulcsár, and H.-F. Raynaud, "Hybrid LQG/integrator control for the VLT eXtreme AO system SPHERE," in *IEEE International Conference on Control Applications, Yokohama, Japan, September 8-10, 2010*.
- [30] U. Frisch, *Turbulence: the legacy of A.N. Kolmogorov*. Cambridge University Press, 1995.
- [31] J. W. Hardy, *Adaptive optics for astronomical telescopes*. Oxford University Press, USA, 1998.
- [32] A. Wirth, J. Navetta, D. Looze, S. Hippler, A. Glindemann, and D. Hamilton, "Real-time modal control implementation for adaptive optics," *Applied optics*, vol. 37, no. 21, pp. 4586–4597, 1998.
- [33] R. Foy and F. C. Foy, "Optics in astrophysics," in *Proceedings of the NATO Advanced Study Institute on Optics in Astrophysics, Cargèse, France 16-28 September 2002*, ser. Nato Science Series II, vol. 198. Springer, 2005.
- [34] Mikhail V. Konnik and José De Doná, "Waffle mode mitigation in adaptive optics systems: a constrained receding horizon control approach," in *Proceedings of American Control Conference (ACC), Washington, DC, USA, June 17–19, 2013*, pp. 3396–3402.
- [35] D. Gavel, "Suppressing anomalous localized waffle behavior in least-squares wavefront reconstructors," in *Proceedings of SPIE*, vol. 4839, 2003, p. 972.
- [36] "QPC - quadratic programming in C," website: <http://sigpromu.org/quadprog>.
- [37] H. Ferreau, H. Bock, and M. Diehl, "An online active set strategy to overcome the limitations of explicit MPC," *International Journal of Robust and Nonlinear Control*, vol. 18, no. 8, pp. 816–830, 2008.
- [38] J. Mattingley and S. Boyd, "Cvxgen: a code generator for embedded convex optimization," *Optimization and Engineering*, vol. 13, no. 1, pp. 1–27, 2012.
- [39] P. Frank, M. Healy, and R. Mastro, "A range-space implementation for large quadratic programs with small active sets," *Journal of optimization theory and applications*, vol. 69, no. 1, pp. 109–127, 1991.
- [40] J. Barzilai and J. Borwein, "Two-point step size gradient methods," *IMA Journal of Numerical Analysis*, vol. 8, no. 1, pp. 141–148, 1988.
- [41] J. Moré and G. Toraldo, "On the solution of large quadratic programming problems with bound constraints," *SIAM Journal on Optimization*, vol. 1, no. 1, pp. 93–113, 1991.
- [42] M. V. Konnik, "Online constrained receding horizon control for astronomical adaptive optics," Ph.D. dissertation, School of Electrical Engineering and Computer Science, the University of Newcastle, Australia, December 2013.
- [43] Graham Clifford Goodwin, Maria Marta Seron, and José De Doná, *Constrained Control and Estimation: An Optimisation Approach*. Springer, Berlin, 2005.
- [44] S. J. Benson, L. C. McInnes, and J. J. Moré, "A case study in the performance and scalability of optimization algorithms," *ACM Transactions on Mathematical Software (TOMS)*, vol. 27, no. 3, pp. 361–376, 2001.
- [45] M. Raydan, "The Barzilai and Borwein gradient method for the large scale unconstrained minimization problem," *SIAM Journal on Optimization*, vol. 7, no. 1, pp. 26–33, 1997.
- [46] D. Goldfarb and A. Idnani, "A numerically stable dual method for solving strictly convex quadratic programs," *Math. Programm.*, vol. 27, pp. 1–33, 1983.
- [47] M. J. D. Powell, "On the quadratic programming algorithm of Goldfarb and Idnani," *Math. Programm. Study*, vol. 25, pp. 46–61, 1985.
- [48] P. Gill, G. Golub, W. Murray, and M. Saunders, "Methods for modifying matrix factorizations," *Mathematics of Computation*, vol. 28, no. 126, pp. 505–535, 1974.
- [49] J. Daniel, W. Gragg, L. Kaufman, and G. Stewart, "Re-orthogonalization and stable algorithms for updating the Gram-Schmidt QR factorization," *Math. Comp.*, vol. 30, no. 136, pp. 772–795, 1976.
- [50] W. Givens, "Computation of plain unitary rotations transforming a general matrix to triangular form," *Journal of the Society for Industrial & Applied Mathematics*, vol. 6, no. 1, pp. 26–50, 1958.
- [51] H. J. Ferreau, P. Ortner, P. Langthaler, L. d. Re, and M. Diehl, "Predictive control of a real-world diesel engine using an extended online active set strategy," *Annual Reviews in Control*, vol. 31, no. 2, pp. 293–301, 2007.

- [52] H. Ferreau *et al.*, “qpoases user’s manual,” *Optimization in Engineering Center (OPTeC) and Department of Electrical Engineering, KU Leuven*, 2011.
- [53] S. Mehrotra, “On the implementation of a primal-dual interior point method,” *SIAM Journal on optimization*, vol. 2, no. 4, pp. 575–601, 1992.
- [54] J. Gondzio, “Multiple centrality corrections in a primal-dual method for linear programming,” *Computational Optimization and Applications*, vol. 6, no. 2, pp. 137–156, 1996.
- [55] L. Vandenberghe, “The cvxopt linear and quadratic cone program solvers,” *Online: <http://cvxopt.org/documentation/coneprog.pdf>*, 2010.
- [56] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, *et al.*, *LAPACK Users’ guide*, SIAM, 1999.
- [57] M. Knapp-Cordes and B. McKeeman, “Improvements to tic and toc functions for measuring absolute elapsed time performance in matlab,” *MATLAB Digest* (doc. 91934v00), Tech. Rep., 2011.
- [58] Y. Nesterov, *Introductory lectures on convex optimization: A basic course*. Springer, 2004.