

Adaptive Optics \mathcal{H}_2 -optimal Control Design Applied on an Experimental Setup

Karel Hinnen^a, Michel Verhaegen^a, Niek Doelman^b

^aDelft Center for Systems and Control, Mekelweg 2, 2628 Delft, The Netherlands;

^bTNO Science and Industry, PO Box 155, 2600 AD Delft, The Netherlands

ABSTRACT

Most adaptive optics systems (AO) are based on a simple control law that is unable to account for the temporal evolution of the wavefront. In this paper, a recently proposed data-driven \mathcal{H}_2 -optimal control approach is demonstrated on an AO laboratory setup. The proposed control approach does not assume any form of decoupling and can therefore exploit the spatio-temporal correlation in the wavefront. The performance of the optimal control approach is compared with a conventional method. An analysis of the dominant error sources shows that the optimal control approach leads to a significant reduction in the temporal error. Since the temporal error grows with the Greenwood to sampling frequency ratio, the performance gain is especially large at large ratios.

Keywords: Adaptive optics, \mathcal{H}_2 -optimal control, spatial-temporal correlation, data-driven disturbance modeling

1. INTRODUCTION

Adaptive optics (AO) systems are often based on a control law that is not able to exploit the spatio-temporal correlation in the wavefront. The most commonly applied control law^{1,2} consists of the cascade of a static part, concerned with the problem of finding the actuator inputs that provide the best fit to the reconstructed wavefront, and a series of parallel feedback loops responsible for stability and closed-loop performance. The special structure of this control law amounts to the implicit assumption that the spatial and temporal dynamics can be decoupled. Even though the extent to which the Taylor frozen flow hypothesis holds might be questionable, this assumption is not very realistic. In general, there exists a strong correlation between the spatial and temporal dynamics of the wavefront, which may be used at the benefit of the controller. Important error sources directly influenced by the AO control system design include the error as a result of measurement noise and the temporal error caused by the inability of the AO system to immediately respond to changes in the wavefront sensor (WFS) measurements. By including a priori knowledge on the spatio-temporal correlation of the wavefront, WFS measurements from the past and all surrounding channels may be used to anticipate future wavefront distortions. In this way, it should be possible to reduce the effect of the delayed response associated with the temporal error. Also the sensitivity to measurement noise may be reduced in this way. If the measurement noise and the temporal error are the dominant error sources in an AO system, the system may therefore benefit from a control strategy that is able to account for the spatio-temporal correlation in the wavefront. Such a control strategy may lead to an improved overall performance, either in terms of the ability to suppress the incoming wavefront distortions, or in terms of the limiting magnitude of the guide star needed for the observations.

In order to improve the performance of the current generation of AO systems, we have recently proposed^{3,4} a data-driven \mathcal{H}_2 -optimal control design strategy that is able to exploit the spatio-temporal correlation in the wavefront. The proposed control strategy consists of two steps. In the first step, a dedicated subspace-identification algorithm is used to identify a full multi-variable atmospheric disturbance model on the basis of open-loop WFS data. The subspace algorithm is sufficiently efficient to identify an atmospheric disturbance model for moderate sized AO systems, without assuming any form of decoupling. In the second step, the identified atmospheric disturbance model is used to compute the optimal controller by formulating the AO control problem in an \mathcal{H}_2 -optimal control framework. In line with our previous work,⁴ it can be shown that if the only dynamics in the

Further author information: (Send correspondence to K. Hinnen)
e-mail: k.j.g.hinnen@tudelft.nl, Telephone: +31 (0)15 2786707

transfer function from the control input to the WFS output is an integer samples delay, then the \mathcal{H}_2 -optimal controller can be computed analytically. Together with the subspace-identification algorithm, this results in a non-iterative and efficient way to go from open-loop WFS data to closed-loop controller design. Since the proposed control design strategy does not assume any form of decoupling, it has the potential to exploit the spatio-temporal correlation imposed by the Taylor hypothesis, without being dependent on it. The goal of this paper is to demonstrate the data-driven \mathcal{H}_2 -optimal control approach on an laboratory setup. The performance of the optimal control approach will be compared with a control law that is often used in the current generation of AO systems. By analyzing the different error sources in the AO system, it will be shown that optimal control is indeed able to reduce the effect of the temporal error. Furthermore, the error analysis provides more insight in the conditions under which the optimal control is to be preferred over the common AO control approach.

The remainder of this paper is organized as follows. Section 2 provides a brief description of the AO laboratory setup used to validate the proposed optimal control approach. In Section 3, we will have a closer look at the problem of modeling the different components in the experimental setup. It will be shown that if the active mirror can be considered static, and the signals are synchronized properly, the transfer function from control input to WFS outputs can be modeled as an integer number of samples delay. Using this knowledge, a data-driven identification is used for modeling the AO system. Section 4 gives a brief summary of the optimal control approach that is validated in this paper. The performance of the proposed optimal control approach will be compared with a conventional AO control law, which will be briefly reviewed in Section 5. This section also provides an overview of the criteria used for performance evaluation. The outcome of the different experiments is described in Section 6. The paper concludes with a short discussion in Section 7.

2. THE AO LABORATORY SETUP

This section considers the AO laboratory setup used to test the proposed control approach. Figure 1 provides a schematic representation of the layout of the optical test bench. In the setup, the objective O1 focuses light from a HeNe laser ($\lambda = 633\text{nm}$) on a pinhole P1. The positive lens L1 creates a collimated beam in order to mimic a distant point source. The wavefront distortions are introduced by the atmospheric turbulence simulator TS. The distorted light beam passes the lenses L2 and L3 that are aligned such that the tip-tilt mirror (TT) is conjugated to the entrance pupil P2.

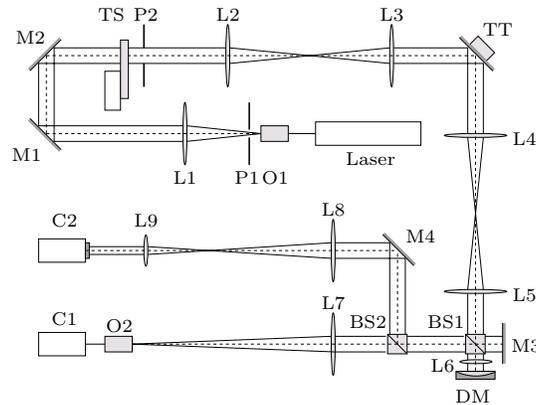


Figure 1. Schematic representation of the optical layout of the AO laboratory setup.

The turbulence simulator consists of a circular plan parallel glass plate that is rotated through the light beam by means of motor drive unit. One side of the glass plate has been machined such that the resulting wavefront distortions have a spatial Kolmogorov distribution. The distortions are characterized by a turbulence coherence length, or Fried parameter,¹ of $r_0 = 2\text{mm}$. With the 10mm aperture P2 representing the entrance pupil of the telescope system, this gives rise to $D/r_0 = 5$, where D denotes the pupil diameter of the system. The rotating glass plate results in a single frozen layer of turbulence satisfying the Taylor hypothesis. Adjustment of the rotational speeds simulates different wind speeds. The distorted light beam passes the lenses L2 and L3 that are aligned such that the tip-tilt mirror (TT) is conjugated to the entrance pupil P2. Separate tip-tilt compensation is important since compensation of this mode by the deformable mirror (DM) would demand too much of its

dynamic range. Via the beam-splitter cube BS1, the lenses L4 and L5 provide an full-size image of the entrance pupil on both the DM and the calibration mirror M3. During normal operation the mirror M3 is shielded; it is only used for calibrating the wavefront sensor (WFS). The DM is a 37-channel electrostatic membrane mirror provided by OKO technologies. The mirror has a circular shaped membrane, 15mm in diameter, and the electrostatic actuators are arranged in a hexagonal grid with an inter-actuator spacing of 0.8mm. A disadvantage of electrostatic actuation is that the actuators are only able to apply a pulling force on the membrane. To allow bi-directional actuation, a bias voltage is applied to each of the actuators. This bias introduces additional focus that is compensated by the negative lens L6 in front of the DM. The second beam-splitter cube BS2 divides the light reflected from the DM in a science path and a WFS path. The camera C1 in the science path, is used to visualize the corrected image. The WFS path consists of a Shack-Hartmann sensor to probe the residual phase errors. The Shack-Hartmann WFS uses a hexagonal array of 127 micro-lenses with a focal distance of 15mm and a pitch of $300\mu\text{m}$. The lenses L8 and L9 reduce the beam size to 3.3mm and ensure that the microlens array is conjugated to the pupil P2. The spot pattern formed by the microlenses is imaged on the WFS camera C2.

The control computer is a general purpose PC with a 3GHz Intel Pentium IV processor, running real-time Linux. This provides a flexible environment for developing, implementing and testing different kinds of control algorithms. Apart from implementing the control algorithm, the control computer is also responsible for processing the frames from the WFS camera. The control computer has to estimate the positions of the spots formed by the microlens array. The spot positions are computed using a standard center-of-mass type of algorithm with background compensation and an adjustable threshold level.^{5,6} Both the centroid algorithm and the control algorithm are implemented in the form of a C-routine, which is executed for each new WFS frame. A second general purpose PC is used for simultaneous recording the images from the science camera C1 and the WFS camera C2. The images recorded by the second PC are only used for performance evaluation.

3. MODELING THE AO SYSTEM

The proposed optimal control approach requires an accurate model of the AO system and the atmospheric turbulence. In this section we consider the problem determining these models. It will be assumed that the wavefront distortion profile ϕ can be represented by a finite-dimensional vector signal $\phi(\cdot) \in \mathbb{R}^{m_\phi}$. A similar representation will be used for the phase correction introduced by the DM and TT-mirror $\phi_m(\cdot)$ and the residual wavefront error $\epsilon(\cdot)$. Whether the vector signals, provide a zonal or modal representation of the wavefront is irrelevant as long as the mean square error of the vector representation provides a good approximation of the mean-square error over the aperture. The proposed \mathcal{H}_2 -optimal control approach is formulated in a discrete-time setting and requires discrete models of the AO system. Since the underlying system evolves in continuous-time, both discrete and continuous-time signals have to be considered in deriving the model structure of the AO components. To distinguish continuous from discrete-time signals the following notational convention will be used. Continuous-time signals have a continuous argument, for instance the continuous version of the phase signal $\phi(\cdot)$ will be denoted as $\phi(t)$, where $t \in \mathbb{R}$. On the other hand, the discrete-time signal obtained by sampling $\phi(t)$ at the time instances kT , $k \in \mathbb{N}$ will be denoted as $\phi(k)$, where k is the discrete-time parameter and T denotes the sampling interval. If no argument is specified, the difference between continuous and discrete-time should be clear from the context. A similar convention will be used to distinguish the Laplace domain from the z-domain. In this way, the Laplace transform of $\phi(t)$ will be denoted as $\phi(s)$ and the z-transform of $\phi(k)$ as $\phi(z)$.

3.1. WFS model structure

An important complication in the AO control problem is that the WFS is not able to directly measure the wavefront distortions $\phi(t)$. Instead of this, the WFS provides some discrete-time signal $s(k)$ that is a measure of the slope or curvature of the wavefront. In this paper we will restrict our attention to the Shack-Hartmann type of WFS, but other types of WFSs can be described in a similar way. Figure 2 shows the block-scheme used for modeling the WFS. The solid lines correspond to continuous-time signals, while the dashed line represent a discrete-time signal. In modeling the WFS, the optical transformation from phase $\phi(k)$ to slopes, or CCD spot positions, is described by a static mapping G ; the so called phase-to-slope geometry matrix. The spot positions on the CCD camera however, can not be observed directly. The CCD camera does not provide instantaneous samples from the spot pattern, but integrates the image over the exposure time $t_e \in \mathbb{R}$. At the end of the

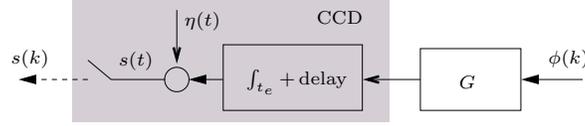


Figure 2. Schematic representation of WFS model

integration, the image is read from the CCD camera, which is then again reset to zero. Furthermore, the time required to read the CCD as well as the processing of the frames, introduces a delay $t_d \in \mathbb{R}$. In Figure 2, the dashed block models the dynamics introduced by the CCD camera. Even though in reality the CCD camera immediately provides a discrete-time signal, the sampling process is modeled separately. The measurement noise introduced by the CCD camera is represented by the additive zero-mean white noise signal $\eta(t)$. In the Laplace domain, the time delay t_d and the integrating action of the CCD camera give rise to the following relation

$$s = \mathcal{G}(s)\phi + \eta \quad \text{where} \quad \mathcal{G}(s) = \frac{1 - e^{-st_e}}{s} e^{-st_d} G = g(s)G. \quad (1)$$

Since the delay and the integrating action affect each of the channels of the WFS in a similar way, the dynamics of the transfer function $G(s)$ is fully decoupled. Furthermore it is clear that the scalar-dynamics $g(s)$ act as a time-invariant low-pass filter. The WFS is used to obtain quantitative information about a wavefront. In general, it is not possible to reconstruct the entire wavefront from the WFS signal $s(t)$. From equation (1), it is clear that only the part of the wavefront that is in the column space of G can contribute to $s(t)$. To arrive at a well-posed control problem, the proposed control approach is based on an alternative representation of the WFS model. This representation considers only the observable part of the wavefront. The reduced representation is obtained by considering the singular value decomposition (SVD)

$$G = U\Sigma V^T = U_1\Sigma_1 V_1^T \quad (2)$$

where $U \in \mathbb{R}^{m_s \times m_s}$ and $V \in \mathbb{R}^{m_\phi \times m_\phi}$ are orthonormal matrices and the partitioning of Σ is such that $\Sigma_1 \in \mathbb{R}^{m_y \times m_y}$ contains all nonzero singular values. Substituting the SVD in equation (1) and pre-multiplying both sides with U_1^T gives rise to the following reduced WFS model

$$y = \Sigma_1 \varphi + \nu \quad (3)$$

where $y(s) \doteq U_1^T s(s) \in \mathbb{R}^{m_y}$, $\varphi(s) \doteq g(s)V_1^T \phi(s) \in \mathbb{R}^{m_y}$ and $\nu(s) = U_1^T \eta(s)$. It can be easily verified that the signal $\varphi(t)$ can be interpreted as a reduced representation of the observable part of the signal obtained by filtering the wavefront $\phi(t)$ with the WFS dynamics $g(s)$. Moreover, both signals have the same Euclidean norm. The signal $y(t)$, on the other hand, can be regarded as a reduced representation of the informative part of $s(t)$. In accordance with the above definitions, the reduced representation of the combined DM and TT phase correction will be defined as $\varphi_m(s) \doteq g(s)V_1^T \phi_m(s)$, while the corresponding residual phase error will be defined as $\varepsilon(t) \doteq \varphi(t) - \varphi_m(t)$. Since the WFS is linear in its input, the corresponding output signal $\varepsilon(t)$ can be expressed as $r(t) \doteq y(t) - y_m(t)$, where $y(t)$ and $y_m(t)$ denote the contributions due to $\varphi(t)$ and $\varphi_m(t)$, respectively.

3.2. Mirror model and discretization of WFS dynamics

In the proposed \mathcal{H}_2 -optimal control approach it is assumed that the transfer function from control input $u(k)$ to WFS output $y(k)$ can be described by an LTI system. As described in Section 6, the electrostatically actuated membrane mirror is almost linear in the applied voltage squared. The static non-linearity can therefore be easily removed by taking the control input equal to the squared voltage applied to each of the actuators. Furthermore, in order to allow bi-directional actuation, a bias voltage is added to the actuators. The bias is chosen in such a way that the initial shape of the DM compensates the additional negative bias lens in front of the mirror. In the proposed \mathcal{H}_2 -optimal control approach there is no separate loop for controlling the TT-mirror. Controlling the TT-mirror is seen as an integral part of the control design problem and the DM and TT-mirror are described by a single integral model. The DM and TT-mirror will both be denoted by the term active mirrors (AMs).

The WFS model in equation (3) provides a continuous-time model of the relation between $\varphi(t)$ and the corresponding WFS signal $y(t)$. Like the signal $s(t)$, the continuous-time signal $y(t)$ is physically non-existent as the WFS provides only a sampled data output. Furthermore, the considered AO system is operated in closed-loop. The proposed \mathcal{H}_2 -optimal control approach requires a discrete-time model of the AO system seen by the controller. This model should describe the relation between the discrete-time closed-loop WFS output $r(k)$, the discrete-time control input $u(k)$ and the open-loop atmospheric wavefront distortion $\varphi(t)$. A schematic representation of the relation between these signals is provided in Figure 3. Since the WFS is linear in its input, it is possible to consider the contributions of the signals $\varphi(t)$ and $\varphi_m(t)$ to the WFS output $r(k)$, separately.

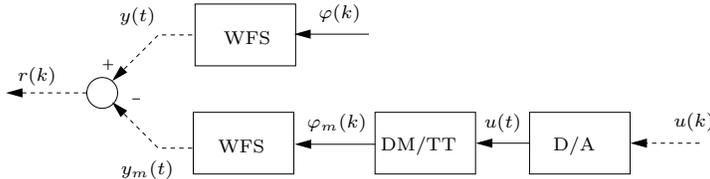


Figure 3. Schematic of AO system as seen by controller

Consider the WFS signal $y_m(k)$ attributed to the wavefront correction $\varphi_m(t)$. The wavefront correction $\varphi_m(t)$, depends on the continuous-time actuator input $u(t) \in \mathbb{R}^{m_s}$. This signal in turn is generated by a digital-to-analog converter D/A with discrete-time control input $u(k) \in \mathbb{R}^{m_s}$. As a result, the contribution of $y_m(k)$ to the closed-loop WFS signal $r(k)$ can be seen as the output of the discrete-time system formed by the cascade of D/A converter, AMs and WFS. Given the continuous-time transfer function of the AMs and the WFS, the equivalent discrete-time transfer function from $u(k)$ to $y_m(k)$ can be computed using the step-invariant transformation.^{7,8} In Looze,⁹ the step-invariant transformation has been used to derive the equivalent discrete-time system for a system with discrete-time measurements based on a CCD camera and an exposure time equal to the sampling interval. The AMs used in the experimental setup have a time constant that is short compared to CCD exposure-time. Hence, from a practical point of view, the mirrors can be considered as static and the only dynamics derives from the D/A converter and the WFS. Even though the corresponding discrete-time transfer function could be derived from the step-invariant transformation, the analysis will be performed in the time-domain as this provides more insight. By pre-multiplying both sides of equation (1) with the matrix U_1^T and using the definition $y(t) = U_1^T s(t)$, the WFS output $y(k)$ can be expressed as

$$y_m(k) = \frac{1}{t_e} \int_{kT-t_e}^{kT} \Sigma_1 V_1^T \phi_m(\tau - t_d) d\tau, \quad (4)$$

where the measurement noise is left out since this will be accounted for in the WFS contribution due to the atmospheric turbulence. Since the AMs are assumed to be static and linear, the projected wavefront $V_1^T \phi_m(t)$ can be expressed as $V_1^T \phi_m(t) = H u(t)$, with $H \in \mathbb{R}^{m_y \times m_y}$ a static influence matrix. The continuous-time control input $u(t)$, is obtained from a ZOH type of D/A-converter. In other words, the discrete-time signal $u(k)$ is converted into a continuous-time signal $u(t)$ by simply holding it constant over the sample interval, i.e. $u(t) \doteq u(k)$ for $kT \leq t < (k+1)T$. Due to the physical limitations of the CCD camera, the exposure time should always be in the range $0 < t_e < T$. According to equation (4), this implies that the output $y_m(k)$ depends at most on two past samples of input $u(k)$. To elaborate the integral, divide the time-delay t_d in an integer number samples delay $d \in \mathbb{N}$ and a remainder $\tau_d \in \mathbb{R}$ as $t_d = dT + \tau_d$, where $d \geq 1$ and $0 \leq \tau_d < T$. Furthermore, let us assume for the moment that $t_e > \tau_d$. Then by substituting the relation $U_1^T \phi_m(t) = H u(t)$ in equation (4) and using the ZOH nature of the input, the WFS output $y_m(k)$ can be expressed as

$$\begin{aligned} y_m(k) &= \Sigma_1 H \left(\frac{1}{t_e} \int_0^{\tau_d} u(k-d) d\tau + \frac{1}{t_e} \int_{T+\tau_d-t_e}^T u(k-d-1) d\tau \right) \\ &= \Sigma_1 H (\alpha_1 u(k-d-1) + \alpha_2 u(k-d-2)), \end{aligned} \quad (5)$$

where $\alpha_1 \doteq \tau_d$ and $\alpha_2 \doteq \tau_d - t_e$ are real-valued coefficients. A similar analysis can be performed for $t_e \leq \tau_d$. In this case the limits of the first integral extends for 0 to t_e , while the contribution due to the second term is zero. As a result, the WFS output $y_m(k)$ can still be expressed as in (6), but with $\alpha_1 = t_e$ and $\alpha_2 = 0$. By using the

reduced WFS model (3) it is clear that for static AMs, the transfer function from $u(k)$ to $\text{varphi}_m(k)$ can be expressed as

$$\mathcal{H}(z) = z^{-d}(\alpha_1 H + \alpha_2 z^{-1} H) = \mathcal{H}(z) = g(z)H, \quad (7)$$

where $g(z) = z^{-d}(\alpha_1 + \alpha_2 z^{-1})$. That this transfer function still includes dynamics, is caused by the particular choice of the reduced wavefront signal $\varphi(t)$. The above derivation shows that for a static mirror, the scalar dynamics $g(z)$ can be expressed as 2 tap finite impulse response plus an integer number of samples delay. In the experimental setup, the real-time software is implemented in such a way that the condition $t_e \leq \tau_d$ is satisfied, with $d = 2$. From the above discussion it is clear that in this case $\alpha_2 = 0$. By absorbing the coefficient α_1 in the influence matrix H , the AMs model reduces to $\mathcal{H}(z) = z^{-2}H$. In this paper, data-driven identification is used to estimate the influence matrix H . To this end, the DM is excited by zero-mean white noise and the resulting WFS response $y_m(k)$ is measured. Considering the reduced AMs model structure, the measured WFS output can be expressed as $y_m(k) = \Sigma_1 H u(k-2)$. From this equation it is clear that an estimate of the influence matrix can be obtained by regressing the signal $\Sigma_1^{-1} y_m(k)$ on $u(k-2)$. Different validation experiments show that the influence matrix identified in this way, provides an accurate description of the measured WFS response $y_m(k)$. The close agreement between the $y_m(k)$ and the WFS signal predicted on the basis of the identified influence matrix H , confirms the choice for the model structure $\mathcal{H}(z) = z^{-2}H$.

Finally consider the WFS contribution $y(k)$ due to the atmospheric wavefront distortion $\varphi(k)$. Since atmospheric turbulence is a continuous-time process, independent from the discrete-time control input $u(k)$, the WFS signal $y(k)$ cannot be characterized in terms of an equivalent discrete-time transfer function. From a control perspective, both WFS contributions $y(k)$ and $y_m(k)$ are therefore entering in a different way. The WFS signal $y(k)$ can be seen as a sampled version of the fictitious continuous-time WFS signal $y(t)$. Assuming that the sampling frequency $f = 1/T$ is sufficiently high, $y(k)$ provides an accurate representation of $y(t)$. Since the WFS model (3) is static, the discrete-time signal $y(k)$ can be immediately related to discretized versions of the signals $\varphi(t)$ and $\nu(t)$. However, seeing the WFS signal $y(k)$ as a sampled version of $y(t)$ does not provide a suitable model for control. In the proposed \mathcal{H}_2 -optimal control approach, it is assumed that the second order statistics of the signal $y(k)$ can be described as the output of an LTI system with a zero-mean white noise input $v(k) \in \mathbb{R}^{m_y}$ and covariance matrix $R_v \doteq \mathcal{E}\{v^T v\}$. This in combination with the static WFS model (3) implies that $y(k)$ and $\varphi(k)$ are modeled as follows

$$\mathcal{S} : \begin{cases} x(k+1) = A_d x(k) + K_d v(k) \\ y(k) = \Sigma_1 C_d x(k) + v(k) \\ \varphi(k) = C_d x(k) + \zeta(k) \end{cases}, \quad (8)$$

where $A_d - K_d \Sigma_1 C_d \in \mathbb{R}^{n_a \times n_a}$ and $A_d \in \mathbb{R}^{n_a \times n_a}$ are stable, and $\zeta(k) = \Sigma_1^{-1}(v(k) - \nu(k))$ is again a zero-mean white noise sequence with covariance matrix $R_\zeta \doteq \mathcal{E}\{\zeta^T \zeta\}$. The atmospheric disturbance model \mathcal{S} will be identified on the basis of open-loop WFS data $y(k)$ and yields a control relevant model of the atmosphere. Since the atmospheric disturbance model is identified from open-loop WFS data, it automatically includes a description of the WFS measurement noise.

4. DATA-DRIVEN OPTIMAL CONTROL FOR AO

This section provides a brief outline of the data-driven \mathcal{H}_2 -optimal control design approach. The approach consists of two major steps, which both have been elaborated in previous work.^{3,4} The first step is concerned with the problem of modeling the turbulent atmosphere. For a disturbance model of the form (8), the problem of modeling $\varphi(k)$ and $y(k)$ reduces to the problem of determining the system matrices A_d , K_d and $\Sigma_1 C_d$. The system matrices are estimated on the basis of open-loop WFS measurement data $y(k)$. Data-driven modeling has the advantage that it provides a good match with the prevalent turbulence conditions and does not depend on restrictive assumptions like the frozen layer hypothesis. Since the model structure does not assume any form of decoupling, it is sufficiently general to describe the spatio-temporal correlation imposed by a frozen flow. A consequence of this rather extensive description, however, is that even for relatively small AO systems the problem of identifying an atmospheric disturbance model is rather challenging. For this reason, a dedicated subspace-identification algorithm has been developed.³ One of the main advantages of the algorithm is that it provides an estimate of the system matrices without the need for spectral factorization.

Given the identified atmospheric disturbance model \mathcal{S} and the combined AMs model $\mathcal{H}(z)$, the second step of the proposed control design approach is concerned with finding the controller $C(z)$ that minimizes the cost-function

$$J = \mathcal{E} \{ \varepsilon^T(k) \varepsilon(k) \} + \mathcal{E} \{ u^T(k) Q u(k) \}, \quad (9)$$

with $Q = Q^T \geq 0$. The regularization matrix Q is used to make a trade-off between the objective of minimizing the expected mean-square residual wavefront error $\mathcal{E}(\varepsilon(k)^T \varepsilon(k))$ and the objective of minimizing the expected amount of control effort $\mathcal{E}(u(k)^T u(k))$. By increasing the regularization matrix Q , it is possible to reduce the amount of energy dissipated by the DM and make the controller more robust to model uncertainties. The matrix Q is typically chosen diagonal, allowing for a penalty on the control effort on each of the actuators separately. The problem of finding the closed-loop optimal controller $C(z)$ can be conveniently expressed in a standard \mathcal{H}_2 -optimal control framework.⁴ Even though straightforward application of \mathcal{H}_2 -optimal control theory provides a very general way of solving the AO control problem, the approach is unable to exploit the special structure in the AO control problem. Since the atmospheric disturbance model with respect to the WFS output $y(k)$ is minimum-phase, the controller can be computed more efficiently. More specifically, if the only dynamics in the transfer function $\mathcal{H}(z)$ from control input $u(k)$ to the applied wavefront correction $\varphi_m(k)$ is an integer number of samples delay, it is possible to derive an analytical expression for the optimal controller. This result is summarized in the following theorem and will be used to compute the optimal controller.

THEOREM 4.1 (OPTIMAL CONTROL WITH QUASI-STATIC DM). *Let the wavefront distortions $\varphi(k)$ and WFS signal $y(k)$ be described by the atmospheric disturbance model (8), with input covariance matrix $R_v > 0$. Furthermore, assume that the only dynamics in the transfer function $\mathcal{H}(z)$ is an integer number of samples delay, i.e. $\mathcal{H}(z) = z^{-d}H$ with $1 \leq d \in \mathbb{N}$. Then, if either Q or H has full column rank, the optimal feedback controller $C(z)$ that is minimizing cost-function (9) is given by:*

$$\begin{bmatrix} \hat{\xi}_1(k+1) \\ u(k) \end{bmatrix} = \begin{bmatrix} \tilde{A} + z^{-d+1}K_d\Sigma_1HF & K_d \\ F(\tilde{A} + z^{-d+1}K_d\Sigma_1HF) & FK_d \end{bmatrix} \begin{bmatrix} \hat{\xi}_1(k) \\ r(k) \end{bmatrix},$$

where \tilde{A} , F and H_Q^\dagger are defined as $\tilde{A} \doteq A - K_d\Sigma_1C_d$, $F \doteq H_Q^\dagger C_d A^{d-1}$ and $H_Q^\dagger \doteq (H^T H + Q)^{-1} H^T$. With a slight abuse of notation, the z -transform parameter is here also used as the unit-shift operator.

The theorem has previously been proved⁴ for the special case of $d = 1$ and $Q = \rho I$, with $\rho \in \mathbb{R}$. The more general case considered here, can be proved by following the same line of reasoning. A full proof of the theorem is beyond the scope of this paper and will be considered in a forthcoming paper. Having an analytical expression for the closed-loop \mathcal{H}_2 -optimal controller is interesting as this leads to an efficient implementation of the proposed control design strategy. The control design strategy obtained by combining the proposed subspace algorithm and the analytical solutions to the \mathcal{H}_2 -optimal control problem, is entirely based on standard matrix operations and provides a non-iterative way to go from open-loop measurement data to closed-loop controller design. The analytical expressions show that the optimal controller can be interpreted as the prediction of future wavefront distortion $\varphi(k)$ over a time horizon of d samples, followed by a static projection on the actuator space.

5. VALIDATION METHODS AND PERFORMANCE MEASURES

The proposed data-driven \mathcal{H}_2 -optimal control approach will be compared with a regularized version of the control approach commonly used in AO. For conformity with the rest of the paper, the approach will be reviewed in terms of the reduced signals. The common AO control approach consists of a cascade of a static matrix multiplication and a series of parallel feedback loops.² Given a new WFS measurement $y(k)$, the static part deals with the problem of finding the DM actuator inputs $u(k)$ that provide the best fit to the wavefront. Let the static relation between $u(k)$ and $y(k)$ be given by $u(k) = Ry(k)$ and let the DM be modeled as $\varphi_m(k) = Hu(k)$. Then, with the static WFS model (3), the problem of finding the reconstruction matrix R is formulated as

$$R = \arg \min_R \left(\mathcal{E} \{ [\varphi(k) - HRy(k)]^T [\varphi(k) - HRy(k)] \} + \mathcal{E} \{ u(k) Q u(k) \} \right), \quad (10)$$

where \mathcal{E} denotes the conditional expectation. In comparison with the usual minimum variance formulation of the reconstruction problem, the above optimization problem includes an additional penalty on the control effort. The

additional regularization is needed to enable a fair comparison between both control approaches. The dynamic range of the DM is small compared to the wavefront distortions introduced by the turbulence simulator. Without input regularization, the actuators of the DM easily saturate, which may even provoke closed-loop instabilities. By noting that the first term in equation (10) can be interpreted as the residual fitting error, it is clear that there is a one to one correspondence with cost-function (9). However, a disadvantage of penalizing the control effort is that it also leads to an increase in the fitting error. Under the assumption that wavefront $\varphi(k)$ and the measurement noise $\nu(k)$ are uncorrelated zero-mean stochastic processes with a Gaussian distribution, the maximum a posteriori estimate of R is given by

$$R = \underbrace{(H^T H + Q)^{-1} H^T}_F \underbrace{C_\varphi \Sigma_1 (\Sigma_1 C_\varphi \Sigma_1 + C_\nu)^{-1}}_E, \quad (11)$$

where $C_\varphi \doteq \mathcal{E}\{\varphi(k)\varphi^T(k)\}$ and $C_\nu \doteq \mathcal{E}\{\nu(k)\nu^T(k)\}$. In the above equation, the operator E provides a minimum variance estimate of the wavefront, while the operator F is responsible for projecting this estimate of the actuator space. Note that the fitting operator F is equal to the projection H_Q^\dagger in the optimal control approach. Since the AO system is operated in closed-loop, the signal obtained from the static reconstruction $u(k) = Ry(k)$ provides only an estimate of the increment needed to the current actuator commands. In order to assure, stability and closed-loop performance, the parallel feedback loops have to possess integrating action. The control law used for performance comparison is given by $u(k) = \beta(1 - \alpha z^{-1})^{-1} Rr(k)$, where α and β are user defined control parameters. As usual, the modification of C_φ , as a result of closed-loop operation, is neglected in computing R .

To obtain a better insight of the conditions under which the proposed data-driven \mathcal{H}_2 -optimal control strategy should be able to outperform the common approach, it is useful to consider the error sources in an AO system. Examples of well-know error sources in AO include anisoplanetic errors, wavefront measurement errors, wavefront fitting errors, temporal errors and tilt related errors.¹ From these error sources, in particular the temporal error is strongly influenced by the control design. Since the temporal error is often one of the major error sources,¹⁰ it is to be expected AO can benefit from an improved control design. In the experimental setup, the total wavefront error is determined almost exclusively by the wavefront fitting error and the temporal error. The wavefront fitting error is the error caused by the fact that the active mirrors cannot take an arbitrary shape. In the case of Kolmogorov turbulence, the mean-square fitting error is given by $\sigma_f^2 = a_f(d/r_0)^{5/3}$, where $a_f \in \mathbb{R}$ is a fitting coefficient depending on the influence functions and d denotes the inter-actuator spacing. With the Fried parameter r_0 of the atmospheric disturbance simulator being fixed, the fitting error can be considered constant. The temporal error on the other hand is caused by the inability of the AO system to immediately respond to changes in the wavefront at the very moment of occurrence. Both bandwidth limitations and pure time delays contribute to this delay in response. Advanced control strategies, like the optimal control approach considered in this paper, should be able to reduce the effect of the time delays by predicting the wavefront distortion at the time of correction. Assume for the moment that the pure time delay is inversely proportional to the sampling frequency and that the bandwidth limitations are proportional to the sample frequency. Furthermore, consider the common control law considered in the previous section, then for Kolmogorov turbulence the mean-square temporal error can be approximated as $\sigma_t^2 = \kappa (f_G/f)^{5/3}$, where $\kappa \in \mathbb{R}$ is a scaling constant and f_G is the so-called Greenwood frequency. The Greenwood frequency can be seen as a characteristic frequency of the atmospheric turbulence. For a single frozen layer with wind velocity v it given by $f_G = 0.427(v/r_0)$ (see¹). The temporal error scales hence as a power law of the Greenwood to sampling frequency ratio (GSFR) f_G/f . Under the assumption that the error sources are uncorrelated, the total wavefront error is given by summing the variances, which gives rise to following error model

$$\sigma_\varepsilon^2 \approx c_0 + c_1 (f_G/f)^{c_2}, \quad (12)$$

where c_0, c_1 and $c_2 \geq 0$ are constants. For the common control approach, the constant c_2 should be close to $5/3$. It will be assumed that the residual wavefront error obtained with other control laws satisfies the same expression. Since the fitting error c_0 is independent from the control design, large performance improvements can only be achieved at large f_G/f ratios. The performance of both control approaches will be compared for different GSFRs. At each GSFR the total residual wavefront error is estimated on the basis of $N_s = 5500$ samples of the residual WFS signal $r(k)$. Given the WFS data $r(k)$, $k \in \{1 \dots N_s\}$, the sample estimate of the mean-square

residual wavefront error is computed as

$$\hat{\sigma}_\epsilon^2 = \frac{1}{N_s m_\phi} \sum_{k=1}^{N_s} \hat{\epsilon}(k)^T \hat{\epsilon}(k), \quad (13)$$

where $\hat{\epsilon}(k) = \Sigma_1^{-1} r(k)$ denotes the reconstructed residual wavefront and m_ϕ is the dimension of unreduced phase representation. For each controller, the computed mean-square residual wavefront errors $\hat{\sigma}_\epsilon^2$ are fitted to equation (12) to obtain an estimate of the coefficients c_0 , c_1 and c_3 . To this end, the constant c_0 is subtracted from the estimates $\hat{\sigma}_\epsilon^2$ and the difference between the logarithm of $\hat{\sigma}_\epsilon^2 - c_0$ and logarithm of $c_1(f_G/f)^{c_2}$ is minimized in a least squares sense. Estimating the coefficients c_0 , c_1 and c_3 is hence formulated as the following optimization problem

$$\min_{c_0, c_1, c_2} \left\| \log_{10}(\bar{\sigma}_\epsilon^2 - c_0) - \mathbb{1} \log_{10}(c_1) - c_2 \log_{10}(\bar{f}) \right\|_2^2, \quad (14)$$

where $\bar{\sigma}_\epsilon^2$ and \bar{f} are the vectors obtained by stacking the different observations of σ_ϵ^2 and the corresponding f_G/f ratios, respectively. Furthermore, $\mathbb{1}$ is a vector of the same dimension as $\bar{\sigma}_\epsilon^2$ with all elements equal to 1. By formulating the fitting problem in this way, the optimization problem becomes linear in the unknowns $\log_{10}(c_1)$ and c_2 . Since for a fixed value of c_0 , the above optimization problem reduces to a standard least squares problem, the optimization problem can be efficiently solved using separable least squares.

Since the ultimate objective of an AO system is to obtain a high resolution image of the science object, it is also interesting to have a number of performance criteria that depend on the optical quality of the corrected image. In the following, the full-width of half maximum (FWHM) and the improvement in peak intensity of the imaged spot will be considered as a measure of performance. To obtain an estimate of these quantities, $N_f = 150$ frames of the science camera have been collected. Subsequently, the obtained images are averaged in order to mimic a long exposure image. Since the science camera has a fixed exposure time of 5ms, the total recording time of the averaged image is constant. If the wavefront would have been perfectly compensated, the image of the point source is diffraction limited resulting an Airy pattern. For this reason the FWHM and peak intensity will be estimated by performing a least squares fit of the theoretical Airy pattern on the time averaged image.

6. EXPERIMENTAL RESULTS

Both the proposed data-driven \mathcal{H}_2 -optimal control approach and the common AO control approach have been implemented on the experimental setup described in Section 2. After aligning the setup, 69 of the 127 microlenses of the Shack-Hartmann WFS are illuminated, resulting in an unreduced WFS signal of 138 channels. The geometry matrix G , specifying the relation between slope measurements and phase, is defined by adapting the well-know Fried configuration for a hexagonal grid. Just as for a rectangular grid, the position of the phase points is determined by shifting the hexagonal grid over half the pitch size. Projecting out the unobservable modes, leads to a reduced WFS signal of 88 independent channels. During the experiments, the WFS exposure time has been adjusted to 5ms. In the experimental setup, the wavefront fitting error is rather large. The DM has only 37 actuators and can therefore only compensate the lowest spatial frequencies of the wavefront. Besides the small number of actuators, also the limited dynamic range and the applied input regularization of the DM is contributing to a large fitting error. Because of the large fitting error, a performance improvement can only be expected at large f_G/f ratios. For this reason, both control algorithms have been tested for GSFRs in the range $f_G/f \in [0.13 - 0.43]$. These f_G/f ratios are indeed large as a f_G/f ratio of 0.5 implies that the characteristic frequency of the turbulence is equal to the Nyquist frequency. The f_G/f ratio is varied by considering different combinations of the GSFR. To find out if the residual wavefront error depends only on the ratio of f_G and f , a number of GSFR combinations with the same f_G/f ratio have been considered.

At each GSFR, the performance of the common control approach is determined considering the same value for the control parameters α and β . The control parameters have been tuned in order to minimize the sample estimate of the mean-square residual wavefront error $\hat{\sigma}_\epsilon^2$ at the median f_G/f ratio. This resulted in the values $\alpha = 0.98$ and $\beta = 0.48$. In evaluating the reconstruction matrix R , the covariance matrix C_φ is computed assuming a perfect Kolmogorov spatial distribution with the Fried parameter r_0 of the turbulence simulator. Furthermore, it is assumed that the measurement noise has a covariance matrix of the form $C_\nu = \sigma_\nu^2 I$, where the variance of the noise σ_ν^2 is estimated from open-loop WFS data $y(k)$ from a static distortion. For a fair comparison, the same

input regularization matrix Q has been used in both control approaches. The regularization matrix is chosen diagonal, i.e. $Q = \text{diag}\{q_1, q_2, \dots, q_{m_u}\}$, and the control effort weighting q_i on each of the actuators is tuned in such a way that actuator saturation is almost avoided and the mirror operates within a linear range. This results in an additional weight on the actuators with a small dynamic range, like the actuators near the edge of the DM. No regularization is needed to keep the TT-mirror within its linear range. In the optimal control approach, the atmospheric disturbance model \mathcal{S} is identified on the basis of $N_s = 5500$ samples of open-loop WFS data $y(k)$. The number of block-rows of past and future measurement data, as used in the subspace-identification algorithm, has been chosen equal to 20. For each GSFR, the model order of the atmospheric disturbance model has been chosen $n_d = 256$.

The outcome of the different experiments is depicted in Figure 4. Each circle and cross in this figure is the result of an experiment at a specified GSFR using the common and optimal control approach, respectively. The ordinate axis of the figure shows the estimated mean-square residual wavefront error $\hat{\sigma}_\varepsilon^2$, while the abscissa axis denotes the GSFR f_G/f . As expected from the discussion in Section 5, the performance improvement is rather

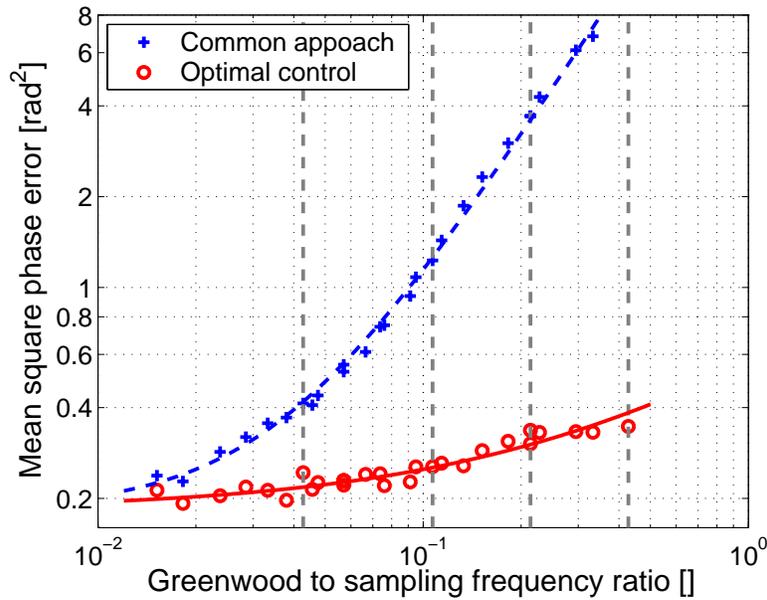


Figure 4. MSE as a function of the Greenwood to sampling frequency ratio f_G/f

low at low GSFRs f_G/f , while a considerable improvement is observed at high f_G/f ratios. The dashed line in Figure 4 represents the fit of the error model (12) to the measurement data $\hat{\sigma}_\varepsilon^2$ obtained with the common control approach. The solid line shows the corresponding fit for the optimal control approach. Since the mean-square residual wavefront error for the optimal control approach shows a rather weak dependence on the GSFR, the parameters c_1 and c_2 are quite sensitive to a slight variation in c_0 . For this reason, optimization over c_0 is omitted and its value is fixed to the value of c_0 found for the common approach. This is a reasonable assumption, since the fitting error in both approaches should be the same. The fitted error model for both control approaches, is given by

$$\hat{\sigma}_\varepsilon^2 \approx 0.1833 + 44.21 \left(\frac{f_G}{f} \right)^{1.6616} \quad (\text{common}) \quad \hat{\sigma}_\varepsilon^2 \approx 0.1833 + 0.3853 \left(\frac{f_G}{f} \right)^{0.7642} \quad (\text{optimal}) \quad (15)$$

Figure 4 shows that the above relations provide an accurate fit to the mean-square residual wavefront errors observed in each of the experiments. The fitted exponent c_2 for the common control approach is close to the theoretical value $5/3$ (≈ 1.667). These observations support the generalized error model proposed in equation (12), which decomposes the total wavefront error in a fitting error independent from the control design and a temporal error influenced by the controller. This implies that at low f_G/f ratios the performance is entirely determined

by the fitting error and nothing can be gained by advanced control. To prove that the constant offset c_0 in the error model can indeed be identified as the fitting error, it is useful to obtain an estimate of the fitting error on the basis of the open-loop WFS data $y(k)$. Starting from the open-loop WFS data collected for identifying the atmospheric disturbance model, the first step in estimating the fitting error is to reconstruct the uncorrected wavefront as $\hat{\varphi}(k) = \Sigma_1^{-1}y(k)$. Recall that both control approaches the fitting operator H_Q^\dagger can be interpreted as the projection of the open-loop wavefront on the actuator space. Neglecting all dynamics, the operator H_Q^\dagger is used in the second step to compute the actuator commands corresponding to the reconstructed wavefront $\hat{\varphi}(k)$ as $\hat{u}(k) = H_Q^\dagger \hat{\varphi}(k)$. The computed actuator commands $\hat{u}(k)$ are then used to determine the wavefront correction $\hat{\varphi}_m(k)$ applied by the active mirrors, from which the fitting error is estimated as the mean-square error of $\hat{\varepsilon}(k) = \hat{\varphi}(k) - \hat{\varphi}_m(k)$. In determining the actuator commands $\hat{u}(k)$ the effect of actuator saturation is investigated by chopping off the input signals if they are out of range. The estimated mean-square fitting error obtained by averaging over the different open-loop data sets is 0.197rad^2 with a standard deviation of 0.007rad^2 . This value is within 2 standard deviations of the constant offset c_0 obtained from the fitted error model (15).

Besides the mean-square residual phase error also the FWHM of the corrected image and an estimate of the gain in peak intensity are used as a performance measure. The different performance measures have been computed for the GSFRs indicated by the vertical dashed lines in Figure 4 and are summarized in Table 1. Apart from these performance measures, the table also includes the ratio of the mean-square residual phase errors $\hat{\sigma}_\varepsilon^2$ and the mismatch error obtained in fitting the observed intensity to the Airy pattern. The results in Table 1 are

f_G/f	Optimal control		Common approach		Relative improvement	
	FWHM	mismatch error	FWHM	mismatch error	gain peak intensity	reduction MSE
$4.3 \cdot 10^{-2}$	14.5px	0.32%	14.5px	0.33%	1.04	1.70
$1.1 \cdot 10^{-1}$	14.5px	0.26%	14.5px	6.14%	1.53	4.82
$2.1 \cdot 10^{-1}$	14.5px	0.28%	27.6px	12.5%	5.70	11.0
$4.3 \cdot 10^{-1}$	14.6px	0.29%	45.1px	6.18%	12.6	20.1

Table 1. Performance comparison control algorithms

consistent with the previous observations. The table shows a considerable performance improvement for large GSFRs, while the performance improvement at low ratios is quite modest due to the relatively large fitting error. Furthermore, the table shows that the optimal control is not only able to achieve a reduction in the observed mean-square residual wavefront error, but is also with respect to optical performance criteria defined in Section 5. The observed image of the point-source has a much smaller FWHM, a larger peak intensity and smaller mismatch error with the theoretical Airy pattern corresponding to a perfect image with any aberrations. A generally used performance measure in AO is the Strehl ratio. Since the gain peak intensity can also be interpreted as the ratio of the Strehl achieved by the optimal and common control approach, it is clear that optimal control is also able to achieve an improvement with respect to this performance measure.

7. CONCLUSIONS

In this paper we have demonstrated a recently proposed data-driven \mathcal{H}_2 -optimal control design strategy^{3,4} on an experimental AO laboratory setup with a frozen flow type of turbulence simulator. In contrast to most existing AO control approaches, the proposed control strategy does not assume any form of decoupling. As a result, it has the potential to exploit the spatio-temporal correlation imposed by the Taylor hypothesis, without being dependent on it. In order to apply the control strategy on the experimental setup, an accurate model of the transfer function from actuator inputs to WFS outputs is required. By analyzing the dynamic behavior of the WFS, it is shown that due to the input-output synchronization in the real-time software and the static nature of the DM and TT-mirror, the discrete-time transfer function from actuator inputs to WFS output reduces to a two samples delay. Considering this model structure, the data-driven identification of the transfer function is formulated as a standard linear regression problem. Furthermore, it can be shown that for this specific case the \mathcal{H}_2 -optimal controller can be computed analytically.

The performance of the proposed data-driven \mathcal{H}_2 -optimal control approach has been compared with a common AO control law. The mean-square residual wavefront error obtained with both control approaches has been estimated over a range of Greenwood to sample frequency ratios (GSFRs). The experiments show that the performance gain achieved by using optimal control is a rapidly increasing function of the GSFR. This observation can be explained by considering the dominant error sources in the AO system. The main contributions to the total mean-square residual wavefront error are the wavefront fitting error and the temporal error. The wavefront fitting error is completely determined by the spatial distribution of the wavefront and the properties of the DM and TT-mirror. Since each of these properties are fixed, the fitting error cannot be influenced by the controller design. On the other hand, the temporal error scales as a power-law of the Greenwood to sampling frequency ratio. By exploiting the spatio-temporal correlation in the wavefront, the optimal controller is able to reduce the effect of temporal error. Since the temporal error becomes the dominant error sources at high GSFRs there is a lot to be gained by optimal control. Even though the measurement noise in the experimental setup is negligible, it may be an error source of considerable importance in real-life telescope systems. Since the optimal control approach is also expected to reduce the effect of the measurement error, it may also be beneficial in situations where the signal to noise ratio is particular low. The optimal control strategy does not only show a performance improvement with respect to mean square residual phase error, but also with respect to the corrected science image. Especially at large GSFRs, optimal control gives rise to a long exposure image with a smaller FWHM, a higher peak intensity and intensity pattern that bears a closer resemblance with an Airy spot.

ACKNOWLEDGMENTS

This research has been conducted in the framework of the Knowledge center for Aperture Synthesis (KAS). The knowledge center is an initiative of TNO TPD, to develop fundamental and advanced technologies for optical aperture synthesis. The knowledge center is a long-term co-operation of TNO and primarily the Delft University of Technology.

REFERENCES

1. J. Hardy, *Adaptive Optics for Astronomical Telescopes*, Oxford series in optical and imaging sciences (Oxford University Press, New York, 1998).
2. F. Roddier, *Adaptive Optics in Astronomy* (Cambridge University Press, 1999).
3. K. Hinnen, M. Verhaegen, and N. Doelman, " \mathcal{H}_2 -optimal Control of an Adaptive Optics System: Part I, Data-driven modeling of the wavefront disturbance," in *Astronomical Adaptive Optics Systems and Applications II*, R. K. Tyson and M. Lloyd-Hart, eds., vol. 5903 of *Proc. SPIE*, pp. 75–85 (San Diego, 2005).
4. K. Hinnen, N. Doelman, and M. Verhaegen, " \mathcal{H}_2 -optimal Control of an Adaptive Optics System: Part II, Closed-loop controller design," in *Astronomical Adaptive Optics Systems and Applications II*, R. K. Tyson and M. Lloyd-Hart, eds., vol. 5903 of *Proc. SPIE*, pp. 86–99 (San Diego, 2005).
5. S. Thomas, "Optimized centroid computing in a Shack-Hartmann sensor," in *ALT'03 International Conference on Advanced Laser Technologies: Biomedical Optics*, R. Wang, J. Hebden, A. Priezzhev, and V.V. Tuchin, eds., pp. 1238–1246 (SPIE, 2004).
6. R. Tyson, *Principles of Adaptive Optics*, 2nd ed. (Academic Press, 1998).
7. K. Aström and B. Wittenmark, *Computer-Controlled Systems, Theory and Design* (Prentice-Hall International Inc., 1997).
8. T. Chen and B. Francis, *Optimal Sampled-Data Control Systems* (Springer-Verlag, London, Berlin, 1995).
9. D. Looze, "Realization of systems with CCD-based measurements," *Automatica* **41**(11), 2005–2009 (2005).
10. M. van Dam, D. L. Mignant, and B. Macintosh, "Performance of Keck Observatory adaptive optics system," *Applied Optics* **43**(29), 5458–5467 (2004).