Minimum variance prediction and control for adaptive optics

Caroline Kulcsár, Henri-François Raynaud, Cyril Petit, Jean-Marc Conan

Minimum variance prediction and control for adaptive optics

Caroline Kulcsár [a], Henri-François Raynaud [a], Cyril Petit [b] and Jean-Marc Conan [b]

[a] *Laboratoire de Traitement et de Transport de l'Information (L2TI), Institut Galilée, Université Paris 13, 93430 Villetaneuse, France*

[b] *Office National d'Études et Recherches Aérospatiales (ONERA), 92320 Châtillon, France*

Abstract

Adaptive Optics (AO) systems enable to compensate the adverse effects of atmospheric turbulence on ground-based telescopes' images in real time, using a deformable mirror (DM) inserted in the telescope's optical path, and measurements provided by a wave-front sensor (WFS). This paper revisits minimum-variance (MV) control design for astronomical AO systems in a state-space framework. It presents a survey of the modeling and control issues arising in this multi-variable disturbance rejection problem. In a linear time-invariant (LTI) framework, and under some mild assumptions, the optimal solution to MV control for AO systems is shown to be a discrete-time LQG controller. This result has been established for a DM with instantaneous response, and for a a fairly general class of DM's dynamics. The state-space approach is extended to Wild-field Adaptive Optics (WfAO) configurations involving several DMs and/or WFSs. Integral-action control used in existing AO systems is compared with LQG controller. Experimental WfAO results obtained on a laboratory test bench are presented, showing significant improvement in performance. Finally, open issues and perspectives of applicative and/or theoretical interests are discussed.

Key words: Adaptive optics, State-space models, MIMO control, LQG optimal control, Kalman filtering, Discrete-time stochastic systems, Sampled data systems

1 Introduction

Nearly sixty years ago, the astronomer Horace W. Babcock made the seminal contribution to what is now called adaptive optics (AO) when he suggested that the angular resolution of ground-based telescopes could be dramatically enhanced by adjusting the surface of a deformable mirror (DM) in response to real-time measurements of the image distortion caused by atmospheric turbulence [1]. Roughly speaking, atmospheric turbulence distorts images because light rays travel at different speeds along different paths. These variations in light rays' paths produce localized phase lags/leads, so that their atmospheric journey results in the addition of a time-varying so-called "turbulent phase" to the image wavefront entering the telescope. Conversely, inserting a DM into the telescope's optical path enables to compensate in real time the turbulent phase with a correction phase, using measurements of the wavefront given by a wavefront sensor (WFS), see Fig. 1. This process yields a (hopefully) smaller residual phase and a corresponding reduction of image distortion.

This work has been partly supported by DGA, Ministère de la Défense, under contract REI N. 0534028, and by CHAPER-SOA ANR-BLAN-0162.

Because the level of DM, sensor and real-time control hardware performance required to implement AO was far beyond the technology available in the 1950s', this brilliantly innovative concept only became a reality – at least in the civilian field – in 1989, with the operational deployment of the COME-ON system on an European Southern Observatory (ESO) telescope [57]. Nowadays, AO technology is routinely employed on all large telescopes, increasing their effective resolution by more than an order of magnitude (for a comprehensive review on AO in astronomy, see, *e.g.,* [56]).

This success inevitably whetted the appetite of the astronomers for AO systems with ever-higher performance, to be integrated in the next generations of ever-bigger and more complex telescopes. Enhancing hardware being an obvious path towards improved performance, significant research efforts have focused on the design of new generations of sensors and actuators with higher spatial resolution and/or better temporal response. New refinements of the original AO concept have also been proposed: XAO for exo-planet detection [18], or wide field AO to achieve the correction needed for the study of stellar populations in nearby galaxies, or active galactic nuclei and first galaxies (with very large red-shifts). Wide field AO is based on the coor-

dinated use of several DMs and WFSs to estimate and correct turbulence in the volume. In the original Multi-Conjugate AO (MCAO) concept [14, 3], several DMs are conjugated in altitude to achieve correction over a wider Field of View (FoV). Since then, numerous variations of the MCAO concept have been proposed (see the articles in [8]); see also [**?**] for a recent survey of advances in AO and their relation to astronomical applications.

Development of control design methods for WfAO is clearly a major challenge, all the more so in the context of Extremely Large Telescopes (ELTs). These future instruments will feature large WFSs and DMs. For example, the European ELT includes an MCAO system with 3 DMs, with diameters up to 2.5m – against 15-20cm in typical current systems. These future DMs will exhibit resonant modes at lower frequencies. They will also feature very large number of actuators and sensors, up to several thousand.

AO performance is ultimately defined as minimizing residual phase variance in the field of interest. Optimizing such performance criterion naturally gives rise to minimum variance (MV) control problems. This paper presents a survey of MV optimal control for AO systems, with special emphasis on several important issues:

- temporal discretization: this includes the delays induced by real-time computers, and also the construction of a discrete-time MV formulation equivalent to the original continuous-time MV problem;
- phase representation spaces and performance analysis on truncated basis;
- analysis of standard AO control used in current on-sky systems;
- extension of the formulation to WfAO, in which priors on turbulent phase play a critical role;

These issues are addressed using a Linear Quadratic Gaussian (LQG) formalism. The time-domain approach to MV control provides a convenient framework for handling the effect of temporal discretization. In the absence of DM's dynamics, it enables the straightforward determination of the discrete time controller minimizing the continuous time MV criterion. The potential performance improvement using LQG control in MCAO is illustrated with an experiment on a laboratory test bench.

Turning these potential improvements into real on-sky performance will require sufficiently adequate yet tractable models of all critical components of the AO loop. These include the AO system *per se* (DM, WFS and real-time controller), but also the disturbance. Constructing and validating such models raise challenging identification for control problems, which are out of the scope of this paper. However, this issue will be dealt with appropriate references throughout this survey.

The organization of the paper is as follows. Section 2 presents an historical overview of optimal AO control. Section 3 introduces the continuous-time AO performance criterion, and shows that for a DM with instantaneous response, the optimal control is obtained as the solution of a tractable discrete-time LQ problem – *i.e.* in reconstructed state-feedback form. Section 4 discusses phase representation on a finite basis, and gives the example of formal computations using Zernike basis on the "inter-sampling variance". Section 5 introduces a general-purpose state-space representation for AO loops, whose critical ingredient is a weak stationary stochastic process modeling the turbulent phase. Section 6 extends the construction of the MV control to DM with linear time invariant (LTI) dynamics. In Section 7, LQG control is compared with the standard integral action AO controller. The state-space approach is extended to WfAO control in Section 8, and experimental results are presented in Section 9 with a test-bench validation in an off-axis AO configuration, a simple case of MCAO. Finally, some conclusions and perspectives are given in Section 10.

2 AO control: problem statement and historical overview

The standard setup of an AO loop is presented in Fig. 1: the correction phase ϕ^{cor} generated by the DM is subtracted from the turbulent phase ϕ^{tur}; the DM control u is computed from real-time measurements y of the residual phase $\phi^{\text{res}} \triangleq \phi^{\text{tur}} - \phi^{\text{cor}}$ provided by the Wave-Front Sensor (WFS). The WFS operates by analyzing the light coming from a guide star, as illustrated in the case of a Shack-Hartmann (SH) WFS in Fig. 2.

Fig. 1. Standard AO setup: the wave-front deformation induced by atmospheric turbulence is compensated by a deformable mirror using closed-loop measurements. Image intensity plots at left illustrate AO correction for a three-star object (simulation).

Fig. 2. Shack-Hartmann Wave-Front Sensor (SH WFS) principle. The light emanating from a guide star located at infinity passes through the atmospheric turbulence, resulting in a distorted incident wave-front. The spot displacement for each sub-aperture is equal to $f \sin \alpha$, where f is the focal length and $\alpha = (\alpha_x, \alpha_y)$ is the average slope.

The SH WFS (see [56], chapter 5) is composed of a grid of sub-apertures optically conjugated with the telescope's entrance pupil, and of a CCD camera located in the focal plane of the sub-apertures grid. An image spot of the guide star is thus formed for each sub-aperture and recorded by the camera. In the absence of turbulence, the incoming wave-front is planar so that all spots are centered. In presence of turbulence, the (x, y) position of each spot shifts according to the local average slope (angles α_x and α_y) of the wave-front. This position is computed *via* appropriate methods (center of gravity or correlation, see, e.g. [**?**]). The SH WFS hence delivers an array of local slope measurements of the wave-front over the whole telescope's pupil. While the SH is common in operational AO systems, some use instead other sensors such as the pyramidal WFS [**?**]. It performs curvature instead of slope analysis, and has therefore a smaller linear variation range.

As atmospheric turbulence evolves rapidly with time, this induces time-varying deformations of the incoming wavefront. From a control point of view, AO control therefore basically boils down to a regulation problem, with the ϕ^{tur} and ϕ^{res} respectively playing the parts of an additive external disturbance and of the error signal (Fig. 3). Control inputs u are for example voltage values in the case of piezo-electric actuators, or squared voltage values in the case of magnetic coils actuators.

This is also a strongly multi-variable control problem with a large number of degrees of freedom: in existing AO systems, u and y are vectors with dozens to hundreds of coordinates for recently upgraded AO systems, e.g. NACO-NAOS [1] at the VLT with 185 actuators and 368 measurements, ALTAIR GEMINI [2] with 179 actuators and 240 measurements or XXXinstru? at Keck Observatory [3] with 349 actuators and 480 measurements. As noted above, this is soon to increase to several thousands for future systems currently designed.

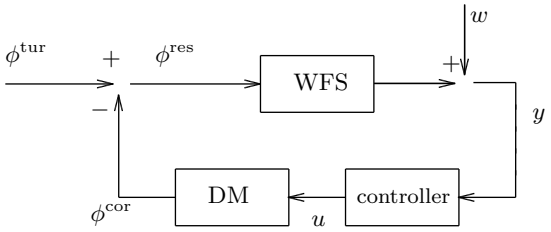

Fig. 3. AO as a disturbance rejection feedback loop, with measurement noise w.

Yet, most AO systems deployed so far rely on a simple control strategy. In essence, starting from the assumption that all blocks in Fig. 3 are linear, the control action combines a static decoupling matrix with a pure 'I' control. In a suitable basis, this corresponds to an integrator with diagonal gain matrix. This type of control, referred to throughout this paper as 'standard AO control', is detailed in Section 7. Quite understandably, efforts to improve control performance have initially focused on improvements of this standard AO control. These include Optimized Modal Gain Integrator (OMGI) where the diagonal gain matrix is tuned according to more refined statistical priors on turbulent phase and measurement noise [21, 22]; alternatively, the pure integrator scalar diagonal controllers can be replaced by higher-order correctors [13, 59], or even by adaptive controllers [23]. However, this does not address a fundamental limitation of the standard approach: because it attempts to diagonalize the closed-loop disturbance rejection transfer function, *i.e.* from ϕ^{tur} to ϕ^{res} in Fig. 3, it rests on the implicit assumption that the DM control, WFS output and residual phase can be made to coincide (at least approximately) through suitable linear transformations.

While this paradigm is tenable in standard AO, it collapses utterly in more complex WfAO configurations. As explained later in Section 8, several WFSs analyze the residual phase in the directions given by some bright guide stars, while DMs conjugated at different altitudes correct the turbulence in the directions given by the objects of interest. Therefore, WFSs no longer measure the residual phase to be minimized. This requires a control approach capable to discriminating between measurement output and variables to be controlled. Much performance is thus to be gained by using state-space control design methods, like H_2 or LQG control. State-space methods also enable to deal with so-called open-loop AO

[1] http://www.eso.org/sci/facilities/paranal/instruments/naco/
[2] http://www.gemini.edu/?q=node/10115
[3] KECK???

systems, where WFS measurements are not affected by control actions. In addition, they enable to incorporate priors on all subsystems of the AO loop (DMs, WFSs, turbulence), and to account easily for additional physical phenomena such as vibrations/windshaking affecting the instrument or DMs' dynamics.

In pioneering works, [45, 44] used a first order DM model together with a continuous time turbulence model of order one for the design of an LQG controller. A very similar formulation was proposed later in [42] using a block-circulant representation diagonalized *via* Fourier transform. However, the problems arising from temporal discretization and pure delays were not properly addressed. In 2002, a discrete time formulation suitable for both classical and MCAO was proposed in [34], with an AR(1) (Auto-Regressive of order 1) turbulence model and a realistic description of control and measurement delays in the loop. This approach was successfully implemented on a real-time laboratory AO bench validating its ability to handle also off-axis AO [49]. An hybrid between LQG and modal control strategies was used in [40, 41], combining a static decoupling matrix with a series of scalar LQG controllers. Recent developments in LQG AO control include optimal control in presence of DM's dynamics [38] or actuator saturation [29], and Fourier domain identification and control [52]. An H_2 AO control with a turbulence model identified using a subspace method is proposed in [25, 26].

A number of authors have also addressed DM related control problems: non-linear control of a piezo-electric actuator, e.g. [61], or control of the DM's surface using PDE models, e.g. [33]. In a different applicative context, high-resolution AO, PDE wavefront models were used to compensate for the non-linearities of a phase-contrast WFS [27].

3 Optimal discrete-time control – fast DM case

In astronomical applications, where image formation involves integrating the optical flow over a large exposure time, a sensible performance criterion is the variance of the residual phase ϕ^{res}, defined as

$$J^{\mathrm{c}}(u) \triangleq \lim_{\tau\to+\infty} \frac{1}{\tau}\int_0^\tau \left\|\phi^{\mathrm{res}}(t)\right\|^2 \mathrm{d}t \qquad (1)$$

where $u \triangleq (u_k)_{k\in\mathbb{N}}$ denotes the sequence of control decisions $u_k \triangleq u(kT)$, which are applied *via* a zero-order hold (ZOH) with a sampling period T. However, it is worth noting that many practitioners in the field of telescope design prefer to use the so-called Strehl ratio (SR). It is defined as the ratio between the central peak of the point spread function (PSF) achieved by the telescope and the central peak of the ideal PSF in the absence of turbulence (the 'Airy spot'). The SR thus varies between 0 and 1 as a percentage, and higher SR means better images (hence better correction). It is shown [**?**] that maximizing the SR amounts to minimizing the residual variance J^c. Furthermore, for small enough values of ϕ^{res} (and appropriate scaling of the phase vectors), the SR is approximately equal to $\exp(-J^{\mathrm{c}})$ [16]. The SR can be computed from actual telescope images, provided that the FoV contains a sufficiently bright point-like light source (in other words, a star).

Assume for the moment that the DM's response is fast (compared with T) and can thus be modeled as a simple gain, so that $\phi^{\mathrm{cor}}(t) = Nu_k$ for all $t \in [kT, (k+1)T)$, where N is the so-called mirror's influence matrix (the columns of N correspond to the influence functions, which describe the deformed surface associated with each actuator [4]). The average value of $\phi^{\mathrm{cor}}(t)$ over $[kT, (k+1)T)$ is thus

$$\phi^{\mathrm{cor}}_{k+1} \triangleq \frac{1}{T}\int_{kT}^{(k+1)T} \phi^{\mathrm{cor}}(t)\,\mathrm{d}t = Nu_k\,. \qquad (2)$$

This relation makes apparent a structural control delay of one frame T. Then, for all t in $[kT, (k+1)T)$, the integrand in (1) can be written as

$$\begin{aligned}\left\|\phi^{\mathrm{res}}(t)\right\|^2 &= \left\|\phi^{\mathrm{tur}}(t) - Nu_k\right\|^2 \\ &= \phi^{\mathrm{tur}}(t)^{\mathrm{t}}\,\phi^{\mathrm{tur}}(t) + u_k^{\mathrm{t}}N^{\mathrm{t}}Nu_k - 2u_k^{\mathrm{t}}N\phi^{\mathrm{tur}}(t)\,.\end{aligned} \qquad (3)$$

Averaging this expression over $[kT, (k+1)T)$ and summing up the results over successive sampling intervals, one gets

$$\begin{aligned}J^{\mathrm{c}}(u) = &\lim_{\tau\to+\infty}\frac{1}{\tau}\int_0^\tau \left\|\phi^{\mathrm{tur}}(t)\right\|^2\mathrm{d}t \\ &+ \lim_{K\to+\infty}\frac{1}{K}\sum_{k=0}^{K-1}\left(u_k^{\mathrm{t}}N^{\mathrm{t}}Nu_k - 2u_k^{\mathrm{t}}N\phi^{\mathrm{tur}}_{k+1}\right)\end{aligned} \qquad (4)$$

where $\phi^{\mathrm{tur}}_{k+1}$ is defined as in (2) as the average of the disturbance over $[kT, (k+1)T)$, *i.e.* $\phi^{\mathrm{tur}}_{k+1} \triangleq \frac{1}{T}\int_{kT}^{(k+1)T}\phi^{\mathrm{tur}}(t)\,\mathrm{d}t$. In the sequel, this notation will be applied to other variables, e.g., x_k will stand for the average of x over $[(k-1)T, kT)$.

[4] AO systems actually often use two physically distinct control devices: a mirror whose surface can be deformed using a grid of piezo-electric or electro-magnetic actuators, and a flat 'tip-tilt' mirror with two degrees of freedom which, as a rule, exhibit markedly different dynamics. This means that the control input is the concatenation of the two control voltage vectors. However, for the sake of simplicity, the combination of the deformable and tip-tilt mirrors shall be referred to throughout this paper, as the DM.

Obviously, the first term in the right-hand side of (4) does not depend on u, so that minimizing $J^{\mathrm{c}}(u)$ is strictly equivalent to minimizing the remaining part of this expression. The part which does depend on u can be transformed into an equivalent *bona fide* discrete LQ criterion by the simple expedient of adding another term independent of u. More precisely, taking

$$\begin{aligned} J^{\mathrm{d}}(u) &\triangleq \lim_{K\to+\infty} \frac{1}{K} \sum_{k=0}^{K-1} \left\| \phi_{k+1}^{\mathrm{tur}} - N u_k \right\|^2 \\ &= \lim_{K\to+\infty} \frac{1}{K} \sum_{k=0}^{K-1} \Big(\left(\phi_{k+1}^{\mathrm{tur}}\right)^{\mathrm{t}} \phi_{k+1}^{\mathrm{tur}} \\ &\quad + u_k^{\mathrm{t}} N^{\mathrm{t}} N u_k - 2 u_k^{\mathrm{t}} N \phi_{k+1}^{\mathrm{tur}} \Big) , \end{aligned} \tag{5}$$

it is apparent that $\arg\min_u J^{\mathrm{c}}(u) = \arg\min_u J^{\mathrm{d}}(u)$. Also, for any choice of u, one gets $J^{\mathrm{c}}(u) \geq J^{\mathrm{d}}(u)$ and

$$\begin{aligned} \varepsilon_{\mathrm{sampl}}^2 &\triangleq J^{\mathrm{c}}(u) - J^{\mathrm{d}}(u) \\ &= \lim_{K\to+\infty} \frac{1}{K} \sum_{k=0}^{K-1} \left(\frac{1}{T} \int_{kT}^{(k+1)T} \left\| \phi^{\mathrm{tur}}(t) - \phi_{k+1}^{\mathrm{tur}} \right\|^2 \mathrm{d}t \right) . \end{aligned} \tag{6}$$

This "inter-sampling variance" is exactly the value of J^{c} that would be achieved if one could somehow manage to make J^{d} equal to zero. Therefore, it should be regarded as an incompressible performance penalty resulting from the use of a discrete-time control with sampling period T. It depends neither on the control law nor on the DM's influence functions. It can be explicitly computed under the assumption that the continuous phase $\phi^{\mathrm{tur}}(t)$ is a stationary ergodic process with known power spectral density [46], see Section 4.

Minimizing J^{d} eventually turns out to be a degenerate LQ problem, since its solution does not require solving a control Riccati equation, but is obtained by simply projecting successive values of $\phi_{k+1}^{\mathrm{tur}}$ onto the mirror space $\mathrm{Im}(N)$. Assume, with no loss of generality (since this property can always been recovered by pruning redundant coordinates of u), that N has full column rank, *i.e.* that $\mathrm{Im}(N) = \dim(u)$, or equivalently that $N^{\mathrm{t}}N$ is invertible. Then, the optimal full information disturbance feedforward control is

$$u_k^{\mathrm{lq}} \triangleq \left(N^{\mathrm{t}}N\right)^{-1} N^{\mathrm{t}} \phi_{k+1}^{\mathrm{tur}} . \tag{7}$$

In practice, the DM is not able to reproduce any turbulence phase: due to the finite number of actuators, its spatial frequencies are limited, whereas turbulent phases have unlimited spatial bandwidth. This produces a so-called "fitting error": would the turbulent phase be perfectly known, this error corresponds to the turbulent deformations that the DM cannot reproduce. This translates into the following lower bound for J^d

$$\varepsilon_{\mathrm{fit}}^2 = J^{\mathrm{d}}(u^{\mathrm{lq}}) \leq J^{\mathrm{d}}(u). \tag{8}$$

The control u_k shall in practice be computed from past controls and WFS measurements $y_0, \ldots, y_k$ available up to time $t = kT$. In this case, the stochastic separation theorem applies: assuming that $\{\phi_k^{\mathrm{tur}}\}$ is a weakly stationary stochastic process, the optimal control then becomes

$$\begin{aligned} u_k^{\mathrm{opt}} &\triangleq \arg\min_{u_k} \mathrm{E}\left(\left\| \phi_{k+1}^{\mathrm{tur}} - N u_k \right\|^2 | \mathcal{I}_k \right) \\ &= \left(N^{\mathrm{t}}N\right)^{-1} N^{\mathrm{t}} \widehat{\phi}_{k+1|k}^{\mathrm{tur}} \end{aligned} \tag{9}$$

where $\widehat{\phi}_{k+1|k}^{\mathrm{tur}}$ is the MV estimate of $\phi_{k+1}^{\mathrm{tur}}$ conditioned to the set $\mathcal{I}_k$ of past measurements (and controls), in other words the conditional expectation $\mathrm{E}\left(\phi_{k+1}^{\mathrm{tur}} | \mathcal{I}_k\right)$. As a result, the original optimal control problem is, for all practical purposes, transformed into an optimal MV prediction problem – the solution of which, in a standard linear Gaussian framework, can be computed using a Kalman filter. So far, it has been implicitly assumed that DM and WFS operate synchronously, in other words that controls are applied and measurements obtained at sampling times $t = kT$. The asynchronous case, resulting in an arbitrary delay between DM and WFS, can be modeled by assuming that measurements are obtained at sampling times $kT - \tau$, while controls are still applied at kT. The optimal solution in both complete and incomplete information cases are given by the same equations (7) and (9). The difference is that the conditional expectation in (9) must be computed with respect to measurements up to time $kT - \tau$. This would be optimally obtained with a different Kalman filter, based on a dynamical model that describes the average phase evolution on intervals of length the greatest common divisor of τ and T if they are commensurate [55], or with a continuous model as in Section 6 otherwise [39**?**]. Asynchronous DM/WFS operation has been accounted for in several papers, e.g. [26, 53]. These approaches, however, are optimal under different and possibly restrictive assumptions on the inter-sampling behavior of the turbulent phase.

4 Phase representation

In order to compute the MV predictor $\widehat{\phi}_{k+1|k}^{\mathrm{tur}}$, a state-space representation of the AO loop shall be needed. This model is required to encapsulate the discrete-time dynamics and mutual interactions of the DM control u, the WFS measurement y, together with a stochastic model of ϕ^{tur}. A first important issue to be addressed is the representation of the phase variables: while u and y are by construction vectors of finite size, ϕ^{tur} is a continuous surface with unlimited spatial bandwidth. Formally

speaking, its domain is a separable Hilbert space, which can therefore be equipped with an infinite countable basis of orthonormal functions.

Throughout this paper, we use a Zernike basis to represent both turbulent and correction phases. The polynomial Zernike basis [43] is widely used in AO because it enables an orthonormal decomposition defined on a support with circular symmetry (e.g. a telescope's pupil), and because the polynomials Z_i have an analytical expression. Transformation to and from a Zernike basis is therefore computationally easy. The decomposition coefficients depend on two subscripts corresponding to radial order n and azimuthal order m, and on polar coordinates $\boldsymbol{r} = (r, \theta)$:

$$\begin{aligned} Z_i &= \sqrt{n+1}\, R_n^m(r)\sqrt{2}\cos(m\theta) \ \text{ for } m \neq 0 \text{ and } i \text{ even}, \\ Z_i &= \sqrt{n+1}\, R_n^m(r)\sqrt{2}\sin(m\theta) \ \text{ for } m \neq 0 \text{ and } i \text{ odd}, \\ Z_i &= \sqrt{n+1}\, R_n^0(r) \ \text{ for } m = 0, \end{aligned} \tag{10}$$

with

$$R_n^m(r) = \sum_{s=0}^{(n-m)/2} \frac{(-1)^s (n-s)!}{s!\left(\frac{n+m}{2}-s\right)!\left(\frac{n-m}{2}-s\right)!} r^{n-2s}. \tag{11}$$

The values of n and m are always integral and satisfy $m \leq n, n - |m| = \text{even}$. The index i is a mode ordering number and is a function of n and m. This gives, for the first modes, the following indexing: for radial order $n = 0$, azimuthal order is $m = 0$ ($i = 1$). For $n = 1$, $m = -1, 1$ ($i = 2, 3$), for $n = 2$, $m = -2, 0, 2$ ($i = 4, 5, 6$) and so on.

Then, for any variable ϕ defined on a disk S,

$$\phi(\boldsymbol{r}) = \sum_{i=1}^{\infty} a_i Z_i(\boldsymbol{r}) \ \text{ with } \ a_i = \frac{1}{S}\int_S \phi(\boldsymbol{r}) Z_i(\boldsymbol{r}) \mathrm{d}\boldsymbol{r} \tag{12}$$

as the basis is orthonormal. High radial orders correspond roughly to high spatial frequencies. Moreover, low radial orders are related to usual optical aberration modes. For example, piston corresponds to first coefficient (radial and azimuthal orders 0), tip and tilt to second and third ones (radial order 1 and azimuthal orders 0 and 1), defocalization and astigmatisms to fourth, fifth and sixth ones (radial order 2 and azimuthal orders 0 to 3). Truncation level of the basis is chosen so as to cut high spatial frequencies that are supposedly negligible. In addition, the piston mode, which corresponds to the average value of the turbulent phase over the telescope's pupil, does not influence performance and does not register on the WFS. It can therefore be removed altogether from the representation.

Note that Zernike modes are a variant of frequency domain representations, allowing to concentrate energy on the low order modes. Similarly to Fourier representations, they are not spatially localized. This is in general a drawback in model-based AO, precisely because both DM controls and WFS measurements involve spatially localized operations. The statistical properties of turbulent Zernike modes have been extensively studied. In particular, their temporal power spectral density (PSD) and spatial cross-correlation are easily computable using physical models [43, 9]. This is convenient to construct control-oriented turbulent models (see Section 5). The availability of Zernike modes' PSD $S_{a_i}(f)$ and the use of an orthogonal basis are also convenient to derive analytical expressions, e.g., for the inter-sampling variance $\varepsilon^2_{\text{sampl}}$ in (6). Indeed, using Parseval's theorem, $\varepsilon^2_{\text{sampl}}$ can be expressed as (see the Annex for detailed calculations)

$$\varepsilon^2_{\text{sampl}} = \text{trace}\left(\int_{-\infty}^{+\infty} S_{\phi^{tur}}(f)\mathrm{d}f - \int_{-\infty}^{+\infty} S_{\psi^{\text{tur}}}(f)\mathrm{d}f\right) \tag{13}$$

where $S_{\phi^{tur}}(f)$ is the continuous-time turbulent phase's PSD and $S_{\psi^{tur}}(f)$ is the PSD of the averaged discrete-time turbulent phase ϕ_k^{tur}:

$$S_{\psi^{tur}}(f) = |\text{sinc}(\pi f T)|^2 S_{\phi^{tur}}(f). \tag{14}$$

This gives

$$\varepsilon^2_{\text{sampl}} = \sum_{i=1}^{\infty}\int_{-\infty}^{+\infty} (1 - |\text{sinc}(\pi f T)|^2) S_{a_i}(f))\mathrm{d}f, \tag{15}$$

where $S_{a_i}(f)$ is the PSD of the Zernike coefficient a_i. Analytical expressions for $S_{a_i}(f)$ are given in [9], thus enabling to evaluate $\varepsilon^2_{\text{sampl}}$ for different turbulence and sampling conditions, see [46].

XXXBlabla bases zonales avec gdes dimensions+ references (Ellerbroek, Gilles, Vogel,+?).

While any control law is ultimately bound to be implemented on a finite basis, one might nevertheless perform controller design in a non finite (*i.e.* PDE-based) framework. Throughout this paper, we choose instead a finite representation. This clearly has the merit of simplicity. The obvious question is whether it may lead to significant performance degradation. In a reconstructed feedforward approach, explicit bounds for performance degradation resulting from truncation can be established on the condition that the basis concentrates the energy of disturbance (turbulent phase) and of control input towards low order coefficients, which is the case for modal basis.

Robustness considerations should also be taken into account, as all effects present in a real system are not

completely modeled. These include for example calibration errors, slight misalignments between DM and WFS, which result in high spatial frequencies. The truncated basis should then be large enough to avoid aliasing of these high order modes on modes associated with lower spatial frequencies. Indeed, the turbulent phase, as a continuous surface with non limited spatial frequency bandwidth, shall be assumed to evolve in a (possibly infinite-dimensional) turbulent space $\mathcal{P}$, which may be different from both the spaces of achievable correction phase (the so-called DM space) and from the space of WFS-observable phase (WFS space). Let us assume, with no loss of generality, that the DM and WFS spaces correspond to closed subspaces of $\mathcal{P}$. Assume further that $\mathcal{P}$ is a separable Hilbert space, and that the turbulence has been approximated on a finite basis, generating a subspace $\mathcal{P}^{\parallel} \in \mathcal{P}$. Denote as $\mathcal{P}^{\perp}$ the orthogonal complement of $\mathcal{P}^{\parallel}$ in $\mathcal{P}$, so that $\mathcal{P} = \mathcal{P}^{\parallel} \oplus \mathcal{P}^{\perp}$. Thus, any turbulent phase can be uniquely decomposed as

$$\phi^{\mathrm{tur}} = \phi^{\mathrm{tur}\parallel} + \phi^{\mathrm{tur}\perp}, \tag{16}$$

with $\phi^{\mathrm{tur}\parallel} \in \mathcal{P}^{\parallel}$ and $\phi^{\mathrm{tur}\perp} \in \mathcal{P}^{\perp}$. The part of the turbulent phase's power that lies in $\mathcal{P}^{\perp}$ is denoted by $\varepsilon^2_{\mathrm{tur}}$.

These modeling errors have an influence on performance in two different manners: through the DM, which shall not reproduce a perfect turbulent phase (this of course includes the fitting error 8), and through the WFS which induces aliasing due to spatial sampling that cannot respect Shannon-Nyquist sampling theorem (remember that the phase has unlimited spatial bandwidth).

The performance degradation due DM's influence functions truncation and DM fitting error can be easily quantified. Whatever the control law, the correction phase $\phi^{\mathrm{cor}} = Nu$ lies in a subspace of $\mathcal{P}$ which may have a non-empty intersection with $\mathcal{P}^{\perp}$, leading to

$$\begin{aligned}\phi^{\mathrm{cor}}_k &= \phi^{\mathrm{cor}\parallel}_k + \phi^{\mathrm{cor}\perp}_k \\ &= N^{\parallel}u_{k-1} + N^{\perp}u_{k-1},\end{aligned} \tag{17}$$

where $N^{\parallel}$ (resp. $N^{\perp}$) corresponds to the projection of influence functions onto $\mathcal{P}^{\parallel}$ (resp. $\mathcal{P}^{\perp}$). Let $\varepsilon^2_{\mathrm{cor}}$ be the power of $N^{\perp}u$, and assume that u has been computed so as to achieve an expected performance

$$J^{\mathrm{d}}_{\mathrm{exp}}(u) \triangleq \lim_{K\to+\infty} \frac{1}{K}\sum_{k=0}^{K-1} \|\phi^{\mathrm{tur}\parallel}_{k+1} - N^{\parallel}u_k\|^2 \tag{18}$$

whereas the actual value of the discrete-time performance criterion is J^{d}. Using (16), (2) and (5), an upper bound of the power of the uncorrected part is obtained as

$$\begin{aligned}J^{\mathrm{d}}(u) - J^{d}_{\mathrm{exp}}(u) &= \lim_{K\to+\infty} \frac{1}{K}\sum_{k=0}^{K-1} \left\|\phi^{\mathrm{tur}\perp}_{k+1} - N^{\perp}u_k\right\|^2 \\ &\leq \varepsilon^2_{\mathrm{tur}} + \varepsilon^2_{\mathrm{cor}}\end{aligned}$$

for zero-mean turbulent phase and control (or $\leq (\varepsilon_{\mathrm{tur}} + \varepsilon_{\mathrm{cor}})^2$ otherwise). Putting all the budget errors together, namely $\varepsilon^2_{\mathrm{sampl}}$ the temporal sampling effect (15), the fitting error $\varepsilon^2_{\mathrm{fit}}$ (8), and the truncation errors $\varepsilon^2_{\mathrm{tur}}$ and $\varepsilon^2_{\mathrm{cor}}$, on obtains upper and lower bounds for the actual performance J^{c}:

$$\varepsilon^2_{\mathrm{fit}}+\varepsilon^2_{\mathrm{sampl}} \leq J^{\mathrm{c}}(u) \leq J^{\mathrm{d}}_{\mathrm{exp}}(u)+\varepsilon^2_{\mathrm{tur}}+\varepsilon^2_{\mathrm{cor}}+\varepsilon^2_{\mathrm{sampl}}. \tag{19}$$

To reduce the effect of the degradations due to basis truncation on control performance J^{d}, the only possible action is to increase the size of $\mathcal{P}^{\parallel}$ in order to have a better representation of ϕ^{tur} and ϕ^{cor}, and thus to lower $\varepsilon^2_{\mathrm{tur}}$ and $\varepsilon^2_{\mathrm{cor}}$. In practice, this guarantees that the deterioration of closed-loop performance caused by truncation can be kept within acceptable (and computable) bounds on the intuitively reasonable condition that the truncated basis is large enough to capture most of power of both the turbulent and correction phases.

Alternatively, it has been proposed to represent/reconstruct the phase in the DM space $\mathcal{P}^{\parallel} = \mathrm{Im}(N)$ as in [40, 58, 37]. This neglects turbulent modes that are not in the DM space but register on the WFS, leading to a possible degradation of the phase reconstruction. It has also been proposed to represent/reconstruct the phase in the WFS space $\mathrm{Ker}(D)^{\perp}$ as in [25]. In this case, this may also degrade reconstruction since the correction phase may be inaccurately represented (this of course depends on the spatial resolution of the WFS). For both representations, limiting phase reconstruction to the DM or the WFS space results in possible performance degradation due to WFS spatial aliasing (i.e. aliasing of turbulent spatial frequencies above the WFS' Shannon-Nyquist spatial frequency). The impact of aliasing on closed-loop performance is, as noted above, strongly dependent on the spatial resolution of the WFS, but also on the control law. Its detailed treatment is beyond the scope of this paper, see for example [51, 17, 54**?**]. Alternatively, reconstructing the turbulent phase in a larger space allows thus to extrapolate the spatial spectrum of the turbulent phase in accordance with statistical priors derived from physical considerations. It also extends naturally to wide-field AO configurations, where the mismatch between DM/WFS space and turbulent phases increases.

5 State-space model and Kalman filter

Let us now examine the observation model in the incomplete information case, where only noisy and delayed discrete measurements are available for control computation. In the case of a Shack-Hartmann WFS, which is used in our experimental set-up, measurements of the first order derivatives of the phase are deduced from CCD camera images of a guide star (GS). This yields local average slopes of the wave-front, sampled on a spatial grid covering the telescope's aperture. To get sufficiently informative and not too noisy measurements, the optical flow is integrated over an exposure time T_e. We suppose here that WFS read-out and slopes computations do not exceed T_e. On the other hand, WFS read-out and slopes computations are relatively time-consuming procedures, the completion of which typically eats up a large fraction of a sampling interval. Therefore, the simplest and commonly used real-time arrangement is to choose the control sampling period T equal to T_e, to assume synchronous operation of DM and WFS, and to perform all read-out, slopes and control computations within one frame $T = T_e$.

Such a configuration can be modeled by postulating that u_k, *i.e.* the control action applied at time $t = kT$, is computed from measurements $y_0, \ldots, y_k$, with y_k depending on ϕ^{res}_{k-1} (Figure 4).

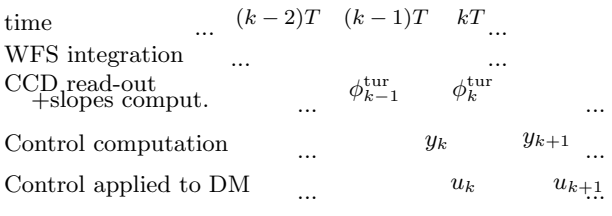

Fig. 4. Chronogram of the AO loop. During time interval (or frame) $[(k-2)T, (k-1)T]$, the wave-front is integrated by the WFS (ϕ^{tur}_{k-1} is the integrated phase) which delivers slopes measurement y_k somewhere in $[(k-1)T, kT]$. The control u_k is then computed to be available at time kT, and is applied on next frame $[kT, (k+1)T]$, leading to a two-frame pure delay in the loop.

Taking into account the one-frame control delay in (2) inherent to the integral nature of the discrete-time AO phase variables, this leads to a total loop delay of $d = 2$ frames. Following the common practice in AO design, the WFS model is assumed to be linear in ϕ^{res} and noisy. As explained above, all computational delays fit into one frame T: measurements y_k shall then be obtained as

$$y_k = D\phi^{\text{res}}_{k-1} + w_k, \tag{20}$$

where w is an additive Gaussian white noise with covariance matrix Σ_w.

Priors on turbulent phase are usually available under the form of a spatial covariance matrix Σ_ϕ with so-called Kolmogorov (or Von Kármán) statistics on Zernike coefficients [28], together with the PSD of each Zernike mode [43, 9]. The PSD is obtained under the so-called Taylor hypothesis, which considers turbulence layers as frozen screens translating at constant wind speed V. Turbulent phase ϕ^{tur} may therefore be assumed to be a Gaussian stationary zero-mean process with known PSD and covariance matrix $\Sigma_\phi \triangleq \mathrm{E}(\phi^{\text{tur}}_k(\phi^{\text{tur}}_k)^{\text{t}})$, where $\mathrm{E}(\cdot)$ stands for mathematical expectation.

Such processes can be approximated by auto-regressive (AR) models. The simplest choice is a vector valued AR(1) model,

$$\phi^{\text{tur}}_{k+1} = A^{\text{tur}}\phi^{\text{tur}}_k + v_k, \tag{21}$$

with v a Gaussian white noise of covariance matrix Σ_v. Taking a diagonal matrix $A^{\text{tur}} = \mathrm{diag}(a_{i,i})$enables to adjust the cut-off frequency for each mode according to priors [**?** 47]:

$$a_{ii} = \exp\left(-\frac{0.3(n(i)+1)VT}{D}\right)$$

where $n(i)$ is the radial order of the i^{th} mode, V the wind speed norm, T the sampling period and D the telescope's diameter. The corresponding autocorrelation function $R_i(m)$ of the i^{th} mode is then $R_i(m) = a_{ii}^{|m|}\Sigma_\phi(i,i)$ where $\Sigma_\phi(i,i)$ refers to element (i,j) of Σ_ϕ. Note that, despite the fact that A^{tur} is diagonal, this construction results in a non diagonal cross-correlation function R_{ij} between Zernike's modes i and j:

$$R_{ij}(m) = a_{ii}^{|m|}\Sigma_\phi(i,j).$$

Once A^{tur} has been determined, the covariance matrix Σ_ϕ is linked to Σ_v through the discrete Lyapunov equation deduced from turbulence model (21): $\Sigma_\phi = A^{\text{tur}}\Sigma_\phi A^{\text{tur}^{\text{t}}} + \Sigma_v$, so that the noise covariance matrix Σ_v is defined as

$$\Sigma_v = \Sigma_\phi - A^{\text{tur}}\Sigma_\phi A^{\text{tur}^{\text{t}}}.$$

The noise covariance Σ_v can thus be chosen so that Σ_ϕ is equal to the spatial covariance of the turbulence phase.

This AR(1) model or its continuous time counterpart has been used by many authors [45, 42, 20, 35, 49, 37, 50]. A better fit of the temporal correlations would require higher order models. Several works have shown, using end-to-end simulations or experiments, that AR(2) turbulent models perform better than AR(1), see for example the precision control of tip-tilt modes for SPHERE [19], for Gemini Planet Imager [**? ?**] or for GeMS, the Gemini MCAO system [**?**]. For global control, AR(2) have also shown their superiority, see [52**? ?**].

As regards numerical values for D and N, they may be obtained by calibrations and/or by using theoretical models (Gaussian influence functions for DM, geometric model for WFS). This is routinely used in AO, with many references available (among others [**? ? ? ? 7?**]).

Let us now turn to the MV estimation problem in this linear Gaussian framework. The first step is to construct $\hat{\phi}^{\rm tur}_{k+1|\mathcal{S}_k}$, the MV estimate of $\phi^{\rm tur}_{k+1}$ based on the sequence of delay- and control-free measurements $\mathcal{S}_k = \{s_0, \ldots, s_k\}$ with $s_k \triangleq y_{k+1}$. The MV estimate $\hat{\phi}^{\rm tur}_{k+1|\mathcal{S}_k}$ is the output of the Kalman filter for the AR(1) model (21) with measurement s_k

$$\hat{\phi}^{\rm tur}_{k+1|\mathcal{S}_k} = A^{\rm tur}\hat{\phi}^{\rm tur}_{k|\mathcal{S}_{k-1}} + L_k\left(s_k - D\hat{\phi}^{\rm tur}_{k|\mathcal{S}_{k-1}}\right), \quad (22)$$

where L_k is the Kalman gain. Since this estimate is to be used for infinite horizon LQG control with stationary models, there is no loss of optimality, with respect to the infinite horizon control criterion J_d, by using the asymptotic Kalman gain [32, 31] given by

$$L_\infty = A^{\rm tur}\Sigma_\infty D^{\rm t}(D\Sigma_\infty D^{\rm t} + \Sigma_w)^{-1}, \quad (23)$$

where Σ_∞ is the solution of the algebraic Riccati equation

$$\begin{aligned}\Sigma_\infty =& A^{\rm tur}\Sigma_\infty A^{\rm tur^t} + \Sigma_v \\ &- A^{\rm tur}\Sigma_\infty D^{\rm t}(D\Sigma_\infty D^{\rm t} + \Sigma_w)^{-1}D\Sigma_\infty A^{\rm tur^t}.\end{aligned} \quad (24)$$

This equation has a unique solution since $A^{\rm tur}$ is stable.

The second step is to construct the predictor $\hat{\phi}^{\rm tur}_{k+1|k}$ for the actual measurement equation (20) with control and measurement delay, *i.e.*

$$y_k = D\phi^{\rm tur}_{k-1} - DNu_{k-2} + w_k. \quad (25)$$

In this case, it is immediately checked that two successive occurrences of $\phi^{\rm tur}$ have to be estimated through

$$\begin{cases} \hat{\phi}^{\rm tur}_{k|k} &= A^{\rm tur}\hat{\phi}^{\rm tur}_{k-1|k-1} \\ &\quad + L_\infty(y_k - D\hat{\phi}^{\rm tur}_{k-1|k-1} + DNu_{k-2}), \\ \hat{\phi}^{\rm tur}_{k+1|k} &= A^{\rm tur}\hat{\phi}^{\rm tur}_{k|k}\,. \end{cases} \quad (26)$$

These equations are simply the non-trivial part of the Kalman filter in predictor form adjusted to the state-space model

$$x_{k+1} = Ax_k + Bu_k + \xi_k \quad (27)$$
$$y_k = Cx_k + w_k \quad (28)$$

where

$$x_k = \begin{pmatrix} \phi^{\rm tur}_k \\ \phi^{\rm tur}_{k-1} \\ u_{k-1} \\ u_{k-2} \end{pmatrix},\ A = \begin{pmatrix} A^{\rm tur} & 0\ 0\ 0 \\ I & 0\ 0\ 0 \\ 0 & 0\ 0\ 0 \\ 0 & 0\ I\ 0 \end{pmatrix},\ B = \begin{pmatrix} 0 \\ 0 \\ I \\ 0 \end{pmatrix},$$
$$\xi_k = \left(v^{\rm t}_k\ 0^{\rm t}\ 0^{\rm t}\ 0^{\rm t}\right)^{\rm t},\ C = (0\ D\ 0\ -DN). \quad (29)$$

Note that the state vector x includes occurrences of u and $\phi^{\rm tur}$. As a consequence, this representation is non-minimal: one could use instead, for example, the smaller state vector $x_k = ((\phi^{\rm res}_k)^{\rm t}, (\phi^{\rm res}_{k-1})^{\rm t}, u^{\rm t}_{k-1})^{\rm t}$. On the other hand, the state transition and control matrices A and B do not depend on the DM/WFS parameters N and D, which makes the model structurally simple and facilitates, among other things, robustness analysis. Also, this choice of state vector allows an easy adaptation to more complex models and configurations. Thus, ARMA turbulence phase models of higher orders can be accommodated by adding additional occurrences of $\phi^{\rm tur}$ and v in the state vector (no additional state is required for an AR(2) model). This choice of state vector enables to easily incorporate prior information on turbulent phase and separated directions of analysis and interest – a key issue in WfAO schemes (see Section 8).

In addition, the state representation can also be regarded as a design tool. In other words, the resulting real-time control should obviously be implemented so as to take advantage of its sparseness and special structure.

6 Optimal control – DM with dynamics

Current AO systems use DMs which are rather small (10 to 30 centimeters), so that their dominant time constants are mostoften neglected as they remain sufficiently small with respect to the sampling period. This may not be true for future telescopes which will include much larger DMs operating at similar sampling rates as e.g. the European Extremely Large Telescope, with a 2.5 m DM and around 1 kHz sampling frequency. Sooner or later, neglecting DM's dynamics in AO control will cease to be an option.

AO loops are sampled-data systems, with controls applied through a ZOH and performance depending on continuous-time variables. In such systems, a critical issue is to control the adverse effects of inter-sampling behavior on performance, see, e.g., [6]. In the special case of AO systems without DM's temporal dynamics, the construction in Section 3 enables to solve the MV problem by breaking the continuous-time criterion in two parts: a discrete-time quadratic criterion J^d which yields the discrete-time MV control, and the control independent

term (inter-sampling variance) which depends only on the turbulence's PSD.

In this case, the identity $\phi^{\mathrm{cor}}(t) = Nu(t)$ needs to be replaced by $\phi^{\mathrm{cor}}(t) = Np(t)$, where $p(t)$ denotes the effective deformation of the DM's surface at time t, and is assumed to be the output of a deterministic continuous-time LTI model with input u.

Using standard results from sampled-date control theory (see, e.g. [12**?** , 5, 6]), this construction of an equivalent discrete-time LQG criterion has been generalized in [38] to the case were a strochastic continuous-time LTI models is available for the turbulence and when the (dsicrete-time) controller is also assumed to be LTI. This approach was subsequently extended to the case of asynchronous measurements, i.e. for arbitrary delay between DM and WFS, enabling in particular to quantify the degradation in performance resulting from neglecting DM dynamics and/or non-integer delays ([38], LOOZE IJC 2010, LOOZE EJC 2011).

A different construction of an equivalent discrete-time LQ criterion, which seeks to exploit the fact that the plant to be controlled is the combination of a stochastic disturbance model and of a deterministic actuator model, has been proposed independently in [10]. This approach was applied to tip-tilt DM for the E-ELT, showing the ability of the optimal control to efficiently compensate for both first- and second-order dynamics (CITER: CARLOS JOSA, CARLOS EJC). It has been also extended to the asynchronous measurements case [**?**] and to woofer-tweeter DM control (CARLOS JOSA).

Both approaches described above can accommodate essentially any deterministic model of the DM's surface deformation. These include in particular matrix-valued second-order equations in the form

$$M_\eta \ddot{\eta}(t) + K_\eta \dot{\eta}(t) + D_\eta \eta(t) = \Psi_e e(t) \qquad (30)$$

$$p(t) = \Psi_p^{\mathrm{t}} \eta(t) \qquad (31)$$

where e denotes the DM's input signal, η is vector of modal coefficients and M_η is an invertible inertia matrix. This class of models can be derived from PDE models of the DM's surface deformation through either direct calculations or FEM approximations.

These deformation models can also be used to force the DM surface p to track the dynamics-free optimal trajectory u^{opt} in section **. This simpler, albeit suboptimal, approach to the AO control problem in presence of DM's dynamics has been pursued in a number of recent publications, using various control techniques [2**? ? ?**]. An important related issue is the identification of sufficiently accurate DM models [**? ? ? ?**].

The minimal set of variables from which the optimal discrete-time control can be computed is then made up of the internal states of the DM's model and of a turbulence model capable of producing ϕ^{tur} and ψ^{tur} as outputs. As an important corollary, any stochastic discrete state-space model capable of providing an adequate description of these statistics and of their relationship to WFS measurements – whether or not they have been derived from underlying continuous-time models – should be regarded as a sufficient model for the MV control of the AO system.

The penalty-free optimal LQG controller, if it is not to result in control actions with intolerable amplitude, should in practice be limited to cases where the DM's dominant time constant remains small compared to the sampling period T. Clearly, less aggressive controls can be obtained by adding an additional penalty on the control energy. In this case, the resulting loss of optimality can be quantified, see Proposition 1 in [**?**].

The case of DM with instantaneous response turns out to be perfectly reasonable for the combinations of real-time computers and actuators used in existing AO systems. This dispenses with the need of solving a control Riccati equation, thereby simplifying the controller design process. Thus, for the sake of simplicity, only this case will be considered from this point down to the end of the paper. However, the construction of the state-space representation and the extension to WfAO in Section 8 are easily adapted to DMs with linear dynamics.

7 Comparison with standard AO control

Since it has been assumed that both DM and WFS are LTI systems, the AO discrete-time optimal control problem can be approached also from an input-output perspective. Indeed, under the customary assumptions that the controller block in Fig. 3 is LTI, that $\{\phi_k^{\mathrm{tur}}\}_{k\in\mathbb{Z}}$ is a weakly stationary stochastic process with PSD $S_{\phi^{\mathrm{tur}}}$, and that $\{w_k\}_{k\in\mathbb{Z}}$ is a white noise independent of ϕ^{tur} with known variance Σ_w, one immediately gets

$$J^{\mathrm{d}}(u) = \frac{1}{2\pi}\int_0^{2\pi} \mathrm{trace}\left(H(e^{j\theta}) S_{\phi^{\mathrm{tur}}}(e^{j\theta}) H(e^{j\theta})^* + H_w(e^{j\theta})\Sigma_w H_w^*(e^{j\theta})\right) \mathrm{d}\theta\,, \qquad (32)$$

where H and H_w are respectively the closed-loop transfer functions from ϕ^{tur} to ϕ^{res} and from w to ϕ^{res}, see Fig.3, and * stands for transpose conjugate. Thus, on the condition that $S_{\phi^{\mathrm{tur}}}$ is known, minimizing J^{d} could be recast as a discrete-time $\mathcal{H}_2$ optimization problem [26]. The solution of this $\mathcal{H}_2$ problem is shown (assuming one is able to construct a stochastically minimal state-space realization of ϕ^{tur}, which can be achieved when its PSD is rational) to be identical to the LQG control law [60]. However, theses two approaches give different frequency/time domain insights in the optimization

problem. Thus, a time-domain framework is considerably more convenient to construct a discrete-time controller minimizing the continuous-time MV criterion. It also enables to solve for example the MV AO problem with actuator saturations in the absence of DM's dynamics [29]. Conversely, an H_2 approach enables to incorporate easily frequency domain priors, and provides a natural connection to H_2/H_∞ or H_∞ formulations [**? ? ?**].

The standard approach to AO control uses a cruder strategy: reduce this multi-variable loop into a series of (quasi) independent scalar ones, and tune them separately using well-honed frequency-domain SISO procedures. The first step is therefore to compute a singular-value decomposition of the so-called interaction matrix DN. This yields two unitary matrices V, W such that $DN = VUW^{\mathrm{t}}$, where U is a diagonal matrix whose general terms U_i are the singular values of DN sorted in descending order. A structural feature of AO systems is the existence of unseen DM modes such as piston and waffle, *i.e.* special DM's shapes which do not register on the WFS. As a result, the interaction matrix necessarily has zero as a multiple singular value, so that U is non-invertible. One can nevertheless define its pseudo-inverse as the diagonal matrix U^+ whose general term is $1/U_i$ if $U_i \neq 0$ and zero otherwise. Using this notation, the standard integral AO controller shall be defined by a temporal domain equation in the form

$$u_k = u_{k-1} + WGU^+V^{\mathrm{t}}y_k\,, \tag{33}$$

where G is a diagonal matrix with general term $g_i \geq 0$ (with $g_i = 0$ when $U_i = 0$), so that the corresponding controller transfer function is $\left(1-z^{-1}\right)^{-1} WGU^+V^{\mathrm{t}}$. Hence, the closed-loop transfer function between the transformed phase variables $\psi^{\mathrm{tur}} \triangleq V^{\mathrm{t}}D\phi^{\mathrm{tur}}$ and $\psi^{\mathrm{res}} \triangleq V^{\mathrm{t}}D\phi^{\mathrm{res}}$ is diagonal, with i-th term equal to

$$H_i\left(z\right) = \frac{1-z^{-1}}{1-z^{-1}+g_iz^{-2}}\,. \tag{34}$$

This scalar disturbance rejection function is in the form $H_i = \left(1+L_i\right)^{-1}$, with $L_i\left(z\right) = g_iz^{-2}/\left(1-z^{-1}\right)$. In the absence of sensor noise, the performance criterion is equal to

$$\begin{aligned} J^d(u) &= \mathrm{E}(\|\psi^{\mathrm{res}}\|^2) \\ &= \sum_i \frac{1}{2\pi}\int_0^{2\pi} \left|H_i(e^{j\theta})\right|^2 S_{\psi^{\mathrm{tur}},i}(e^{j\theta})\mathrm{d}\theta\,, \end{aligned} \tag{35}$$

where $\mathrm{E}(\cdot)$ stands for mathematical expectation. Each term of this sum can be evaluated and/or optimized (by tuning the scalar gains g_i's) independently, using elementary frequency-domain techniques. The standard default setting is $g_i = 0.5$ for all modes, which guarantees closed-loop stability with a gain margin of 3 dB. Alternatively, in the so-called Optimized Modal Gain Integrator approach, the g_i's are tuned independently, using whatever information is available on the PSD of the turbulent eigenmodes defined by the coordinates of ψ^{tur} [21].

Such a standard control approach is not without merits. Some stem from the fact that it is only a multi-variable 'I' control: ability to nullify the impact of constant additive disturbances, such as actuator and/or sensor offsets, plus low real-time computational cost. Another is that it requires only limited *a priori* knowledge on the system, especially since the interaction matrix DN can be experimentally calibrated by the simple (if somewhat tedious) expedient of sending a finely focused laser beam on the DM, pushing each DM's actuator in turn and registering the corresponding effect on the WFS.

How does standard AO control compare with LQG? Let us start with real-time computational burden. In the standard integrator approach, when n_u DM controls are computed from n_y WFS measurements, n_un_y multiplications are required at each time slot. For LQG control estimating n_ϕ coordinates of the turbulent phase and using the AR(1) model (21) with diagonal state transition matrix, back-of-the-envelope calculations show that the minimal real-time computational burden is something in the order of $n_\phi(n_u+2n_y)+n_un_y$ multiplications.

To get a deeper insight, one may reinterpret the standard approach in a state-space framework. It is immediately checked that (33) corresponds to a reconstructed feedforward control in the form (9), where the prediction $\widehat{\phi}^{\mathrm{tur}}_{k+1|k}$ is obtained using a suboptimal observer (*i.e.,* not necessarily a Kalman filter) adapted to an implicit stochastic disturbance model $\phi^{\mathrm{tur}}_{k+1} = \phi^{\mathrm{tur}}_k + \nu_k$ [30**?**]. While this random walk is unstable, and thus inherently prone to induce controller wind-up, it is likely to yield one step-ahead predictions pretty similar to those derived from the slow-moving AR(1) model in Section 5. Note that when integrators are replaced by scalar correctors of higher order [13, 59, 15], this can be reinterpreted as an observer-based control structure adapted to implicit stochastic models of higher order. More precisely, correctors of order n_c would correspond to ARMA(n_c, n_c-1) models of ψ^{tur}.

This state-space reinterpretation of the standard approach reveals an important drawback. The implicit turbulent phase model turns out to be restricted to the subspace $\mathrm{Im}\,(N) \cap \ker\,(D)^{\perp}$ – in plain English, the part of the turbulence which, should it stay constant over time, could be simultaneously analyzed by the WFS and compensated by the DM. This makes sense as long as one only aims at predicting/nullifying future values of y – which is mostly what matters in standard AO, where direction of interest coincides with direction of measurement. This is no longer the case in WfAO, where one

seeks to minimize the variance of the residual phase in direction(s) of interest which happen not to coincide with the direction(s) of analysis, as explained in next section. For such configurations, an LQG control based on a turbulent phase model incorporating relevant prior information on the spatio-temporal correlation structure (including, crucially, the part not seen by the WFS) can be expected to do significantly better.

8 Optimal MCAO control

Current AO systems suffer from anisoplanatism. Due to this effect, the quality of the correction deteriorates dramatically and inexorably as the angular distance from guide stars increases. The root cause of this problem is the thickness of the earth's atmosphere. More precisely, the standard AO set-up in Fig. 1 implicitly assumes that atmospheric turbulence can be reduced to a single flat resulting turbulent layer at telescope's pupil level. Yet, the turbulence actually exhibits a distribution in volume, and should therefore be regarded rather as made up of an infinite number of turbulent layers distributed through the whole depth of the atmosphere, each with a life of its own [11]. The consequence is anisoplanatism: light rays emanating from the guide star and from an object of interest located some angular distance away are affected by different phase shifts. Correction in a direction distinct from the direction of analysis can be then pretty lousy even if the AO control makes the WFS measurement small.

This is where Multi-Conjugate Adaptive Optics kicks in [14, 3]. Opticians have long mastered the trick of conjugating a DM with some arbitrary altitude h so that, while physically located in the telescope's path, it shall generate an equivalent correction wavefront located at altitude h. An AO system featuring several conjugated DMs (associated with several WFSs pointing at different guide stars) should therefore be able to significantly reduce anisoplanatism by compensating the influence of several turbulent layers in all directions [4]. Better still, there is no obvious *a priori* limitation on MCAO performance. Indeed, an ideal telescope fitted with an infinite number of DMs and provided with perfectly accurate advance knowledge of the the three-dimensional distribution of the turbulence could conceivably aim for AO's Holy Grail: perfect compensation in every possible direction.

In practice, the number of DMs, WFSs and turbulent layers in the system's model shall be limited. As noted in the introduction, a number of WfAO concepts have been proposed in recent years. Detailing how the state-space approach could be extended to all those configurations would go beyond the scope of this paper. Instead, for the sake of simplicity, we shall focus here on a prototypical MCAO set-up made up of $n_{\rm M}$ DMs and $n_{\rm GS}$ WFSs pointing at $n_{\rm GS}$ guide stars. Likewise, it shall be assumed that the turbulent phase can be decomposed as $n_{\rm L}$ independent layers. Note that the altitude of turbulent layers and correction screens, nor even their numbers, do not necessarily match.

Consider now a set $\beta \triangleq \{\beta_1, \beta_2, \ldots, \beta_{n_{\rm I}}\}$ of $n_{\rm I}$ directions of interest, each of which is characterized by a pair $\beta_i \triangleq (\beta_{i,1}, \beta_{i,2})$ of angular deviations (typically measured from the center of the field of view). Astronomical objects, being located far away from the atmosphere, are considered to be infinitely far. The cone formed by the light emanating from one object of interest and reaching the telescope's aperture becomes a tube, the diameter of which is equal to the telescope's one, see Fig. 5. The image distortion along direction β_i is encapsulated in the resulting phase obtained by summing the intersections of all turbulent layers and correction screens with this particular tube. This also applies to objects of analysis, *i.e.* guide stars. All turbulent and correction phases need therefore to be defined, for each layer, on a meta-pupil large enough to encompass all tubes of interest and analysis.

Fig. 5. Typical MCAO setup with $n_{\rm M} = 2$ DMs, $n_{\rm L} = 3$ turbulent layers, $n_{\rm I} = 1$ direction of interest β and $n_{\rm GS} = 2$ directions of analysis α_1 and α_2, with $\alpha_2 = 0$.

The turbulent phase $\phi^{\rm tur}$ is now defined as the collection of the turbulent phases $\phi^{{\rm tur},\ell}$, for all turbulent layers $\ell = 1, \ldots, n_{\rm L}$. Likewise, the correction phase $\phi^{\rm cor}$ and the control vector u are respectively the collections of the correction phases $\phi^{{\rm cor},m}$ and of the control tensions u^m for $m = 1, \ldots, n_{\rm M}$. Using this notation, the residual resulting phase in the direction of interest β_i can be

written as

$$\phi^{\rm res}_{\beta_i} \triangleq \sum_{\ell=1}^{n_{\rm L}} M^{{\rm tur},\ell}_{\beta_i} \phi^{{\rm tur},\ell} - \sum_{m=1}^{n_{\rm M}} M^{{\rm cor},m}_{\beta_i} \phi^{{\rm cor},m} \quad (36)$$

$$= M^{\rm tur}_{\beta_i} \phi^{\rm tur} - M^{\rm cor}_{\beta_i} \phi^{\rm cor} \quad (37)$$

where the matrices $M^{\rm tur}_{\beta_i} \triangleq \left(M^{{\rm tur},1}_{\beta_i}, \ldots, M^{{\rm tur},n_{\rm L}}_{\beta_i} \right)$ and $M^{\rm cor}_{\beta_i} \triangleq \left(M^{{\rm cor},1}_{\beta_i}, \ldots, M^{{\rm cor},n_{\rm M}}_{\beta_i} \right)$ are deduced from the MCAO geometry: $M^{{\rm tur},\ell}_{\beta_i}$ and $M^{{\rm cor},m}_{\beta_i}$ are operators which cut out pupil-sized slices of the turbulent layer ℓ and correction screen m in direction β_i; the sum of the contributions of all layers is then performed to obtain the residual residual phase $\phi^{\rm res}_{\beta_i}$ in (36).

A natural performance criterion can now be defined as the sum of the residual phases' variances in all the directions of interest, *i.e.*

$$J^{\rm c}_{\beta}(u) \triangleq \lim_{\tau \to +\infty} \frac{1}{\tau} \int_0^{\tau} \left(\sum_{i=1}^{n_{\rm I}} \left\| \phi^{\rm res}_{\beta_i}(t) \right\|^2 \right) {\rm d}t \, . \quad (38)$$

Assuming as before that each DM has linear and instantaneous response, so that for all t in $[kT, (k+1)T)$, the correction phase of the j-th DM is $\phi^{{\rm cor},j}(t) = N_j u^j_k$, one immediately gets that in this sampling interval

$$\sum_{i=1}^{n_{\rm I}} \left\| \phi^{\rm res}_{\beta_i}(t) \right\|^2 = \left\| M^{\rm tur}_{\beta} \phi^{\rm tur}(t) - M^{\rm cor}_{\beta} N u_k \right\|^2 \quad (39)$$

where N is the so-called generalized influence matrix, defined as a block-diagonal matrix with general term N_j. The linear operators $M^{\rm tur}_{\beta}$ and $M^{\rm cor}_{\beta}$ correspond respectively to the vertical concatenation of $M^{\rm tur}_{\beta_i}$ and $M^{\rm cor}_{\beta_i}$ for $i = 1, \ldots, n_{\rm I}$. They depend only on direction β and altitudes of turbulent and correction layers. Using this notation, and reasoning along the same lines as in Section 3, it is immediately checked that minimizing $J^{\rm c}_{\beta}$ is equivalent to minimizing at each time

$$J^{\rm d}_{\beta,k}(u_k) \triangleq {\rm E} \left(\left\| M^{\rm tur}_{\beta} \phi^{\rm tur}_{k+1} - M^{\rm cor}_{\beta} N u_k \right\|^2 \mid \mathcal{I}_k \right) . \quad (40)$$

Assuming for the sake of simplicity that $M^{\rm cor}_{\beta} N$ has full column rank, this would result in the optimal MCAO control

$$u^{\rm opt}_k \triangleq \arg\min_{u_k} J^{\rm d}_{\beta,k}(u_k)$$
$$= \left((M^{\rm cor}_{\beta} N)^{\rm t} M^{\rm cor}_{\beta} N \right)^{-1} (M^{\rm cor}_{\beta} N)^{\rm t} M^{\rm tur}_{\beta} \hat{\phi}^{\rm tur}_{k+1|k}. \quad (41)$$

Consider now the measurement process. Under the assumption that the output y^j of the j-th WFS is a linear function (with matrix D^j) of the residual phase in the direction α_j, it can be modeled as

$$y^j_k = D^j \left(\frac{1}{T} \int_{(k-2)T}^{(k-1)T} \phi^{\rm res}_{\alpha_j}(t) \, {\rm d}t \right) + w^j_k \quad (42)$$

where w^j is a Gaussian white noise. Denoting as $\alpha \triangleq \{\alpha_1, \alpha_2, \ldots, \alpha_{n_{\rm GS}}\}$ the set of measurement directions, this can be rewritten in compact form as

$$y_k = D \left(M^{\rm tur}_{\alpha} \phi^{\rm tur}_{k-1} - M^{\rm cor}_{\alpha} N u_{k-2} \right) + w_k \quad (43)$$

where D, $M^{\rm tur}_{\alpha}$ and $M^{\rm cor}_{\alpha}$ are constructed similarly to N, $M^{\rm tur}_{\beta}$ and $M^{\rm cor}_{\beta}$.

Finally, let the evolution of each turbulent layer be described for example by an AR(1) model similar to (21), *i.e.*

$$\phi^{{\rm tur},i}_{k+1} = A^{\rm tur}_i \phi^{{\rm tur},i}_k + v^i_k. \quad (44)$$

Denoting as $A^{\rm tur}$ the block-diagonal matrix with general term $A^{\rm tur}_i$, it is obvious that an adequate state representation can be constructed using the same definition of state vector x and of matrices A and B given in (29), with $\phi^{\rm tur}$ being the vertical concatenation of $\{\phi^{{\rm tur},i}\}_{i=1,\ldots,n_{\rm L}}$, and with

$$C \triangleq \begin{pmatrix} 0 & D M^{\rm tur}_{\alpha} & 0 & -D M^{\rm cor}_{\alpha} N \end{pmatrix} . \quad (45)$$

As indicated in Section 5, this straightforward extension from standard AO to MCAO would not be feasible for a choice of state vector including for example occurrences of the residual phase.

The optimal control is obtained, as in the classical AO case, in reconstructed disturbance feedforward form (a disturbance feedforward control where the non measured disturbance is estimated using an observer): the output of the steady-state Kalman filter derived from this state representation is plugged into (41). Note that while this system is structurally neither controllable nor observable, it is nevertheless stable, so that the filtering algebraic Riccati equation has a unique solution which can be computed using standard numerical procedures.

9 Experimental results for an off-axis configuration

Off-axis AO is a particular case of MCAO, where only one on-axis WFS and one DM are considered, with distinct directions of analysis and interest. This configuration has the advantage of making it possible to be tested on any classical AO system, and is also a convenient way to demonstrate anisoplanatism effects. In this section, performance is compared between a standard integrator and LQG control.

The star of interest is distinct from the GS with an angular separation β, and the GS is on-axis, *i.e.* the direction of analysis is $\alpha = 0$, see Fig. 6. A Shack-Hartmann WFS measures on-axis residual phase. A single turbulence layer at altitude h is considered and the unique DM is optically conjugated with the entrance pupil. The light beams from both on-axis and off-axis sources correspond to the two footprints A and B on the turbulent layer, which also correspond to the projections of the telescope's aperture in each direction. Turbulence at altitude h is then predicted in the whole meta-pupil C.

Fig. 6. Off-axis configuration. The GS is on-axis ($\alpha = 0$), whereas the direction of interest is off-axis, the two stars being separated by an angle β. A single layer is considered, at altitude h.

Using notations defined in Section 8, as the DM is conjugated with the telescope's pupil, projectors $M_\alpha^{\rm cor}$ and $M_\beta^{\rm cor}$ are equal to identity. Therefore, criterion (40), measurement equation (43) and optimal control (41) become respectively

$$J_{\beta,k}^{\rm d}(u_k) = \|M_\beta^{\rm tur}\phi_{k+1}^{\rm tur} - Nu_k\|^2, \quad (46)$$

$$y_k = D(M_\alpha^{\rm tur}\phi_{k-1}^{\rm tur} - Nu_{k-2}) + w_k, \quad (47)$$

$$u_k^{\rm opt} = \left(N^{\rm t}N\right)^{-1} N^{\rm t} M_\beta^{\rm tur}\hat{\phi}_{k+1|k}^{\rm tur}. \quad (48)$$

All experiments have been performed on ONERA's test bench BOA (see Fig. 7). This bench includes a turbulence generator (turbulent phase screen), two sources (fibered LASER diodes) and a telescope simulator, the AO system (DM and WFS) and an imaging camera. The turbulence layer is generated thanks to a rotating phase screen mirror (made at Observatoire de Paris-Meudon) which reproduces wind effects and Kolmogorov statistics. The turbulence strength has been chosen weak but in the range values of turbulence in altitude observed at Mount Paranal, Chili (see [8]).

Fig. 7. ONERA test bench BOA: red line indicates the optical path. Dimensions are 2 m×1 m.

Depending on the turbulent layer's altitude h, the intersection of footprints A and B may vary, together with their relative separation defined as $\delta = \beta h/D$, where D is here the telescope's pupil diameter. This is thus a relevant measure of the distance between the sources. A relative separation of δ =40 % roughly corresponds in this set-up to an angular separation of $\beta = 4$ arcmin for a $D = 8$ meter telescope and a turbulent layer at $h = 3000$ m.

Wave-front correction is performed using two mirrors, a tip-tilt mirror with two degrees of freedom, and a 9×9 actuator piezo-electrical DM where only 69 actuators are used. The total number of actuators is thus $n_u = 71$. The Shack-Hartmann WFS has 8×8 sub-apertures, 52 of which are used, leading to $n_y = 104$ coordinates in measurement vector y (52 x-slopes and 52 y-slopes). Turbulence phase is estimated using $n_\phi = 150$ Zernike modes. The real-time computer (a PC under Linux) is running [5] at $T = 60$ Hz. Finally, the imaging camera is a CCD Princeton with 512×512 pixels.

Comparison of different point spread functions (PSFs) obtained experimentally on the imaging camera are presented in Fig. 8 for both stars. At the top (Fig. 8.a),

[5] A new real-time computer built by Shaktiware Inc. since allows 90 estimated modes at 500 Hz or 300 estimated modes at 130 Hz.

no correction is performed (open-loop), and PSFs are simply two blurred spots, giving a Strehl ratio (SR) of 7%. In the middle (Fig. 8.b), on-axis correction is performed thanks to an optimized integrator, and the object of interest (off-axis) is severely blurred. The integrator achieves an SR of 93 % on-axis and performance drops to 34 % in the direction of interest. Of course, off-axis correction is not possible with a standard integrator as it seeks to nullify the measurements. Off-axis optimal LQG control achieves 81 % SR on the object of interest (Fig. 8.c) and 43 % SR on-axis.

Fig. 8. Experimental PSFs for on-axis GS (left column) and off-axis star of interest (right column) a. in open-loop (top), in closed-loop b. with an integrator (middle) and c. with an LQG control (bottom). The relative separation is 20%, and the SRs are respectively from top to bottom and left–right 7%–7%, 93%–34% and 43%–81%.

As relative separation increases, anisoplanatism degrades performance, as illustrated by experimental results reported in Fig. 9: the farthest the object of interest is away from the GS, the least turbulence in the corrected direction is similar to the one analyzed by the WFS. Obviously, the LQG control is a long way ahead of the integrator because correction is performed in the desired direction, with an SR of about 60 % for LQG control at 30 % of relative separation, against only 20 % SR with the integrator.

10 Conclusion and some perspectives

From the vantage point of control theory, the approach to AO control design presented in this paper can be viewed as a special case of discrete-time reconstructed disturbance feedforward control, in a standard, albeit strongly multi-variable, LTI framework. From an applicative point of view, its main attraction lies undoubtedly in the inherent flexibility of the observer-based control structure. Indeed, working along the same lines as in Section 8, the LQG construction can be almost effortlessly extended to the new AO concepts mentioned in the introduction by the essentially simple process of making appropriate modifications to the performance criterion, the state vector and the measurement equation. These include open-loop AO systems, where DMs

Fig. 9. Experimental performances of LQG control compared with integrator, as a function of the relative separation. Error bars are indicated, corresponding to experiments realized on different days.

are positioned in the telescope's path so as not to affect WFS measurements. Such configurations have been proposed for wide-field of view imaging, where a large number of small DMs are to perform local corrections around objects of special interest (multi-object AO, see e.g. [24]). An alternative control strategy in open-loop AO is to perform phase reconstruction based on the last available WFS measurement (an inverse problem), and to project orthogonally the estimate onto the DM, see e.g. [**?**] and references therein, or [**?**]. A variant is to reconstruct the measurements in the direction of interest and to project them onto the DM, which has been applied to the CANARY multi-object AO on-sky demonstrator [**? ?**].

Another potentially interesting extension is vibration filtering. By inserting in the model additional states corresponding to spring-mass subsystems, observer-based control can filter out and/or compensate telescope vibrations. This has been validated in simulations and test-bench experiments [48**? ? ?**]. The reconstructed feedback structure provides a natural anti-wind-up mechanism, and performance degradation due to actuator saturation can be easily quantified in absence of DM dynamics [29].

The state-space approach also authorizes temporal flexibility. More precisely, it is readily checked that all the constructions presented in this paper can be adapted to multi-rate configurations where the WFS and DM operate at different sampling rates [55**?**], and/or in an asynchronous manner [53, 26**? ?**]. This may turn out to be especially relevant in astronomical AO systems, where the WFS's CCDs operate in conditions of weak illumination, so that making them run too fast results in exploding measurement noise. This limitation could conceivably been circumvented by updating the control action several times between two consecutive WFS measurements. However, it should be noted that this would

put more demand on the predictive power of the turbulent phase model, because the Kalman filter would have to operate in "observer open-loop" pure prediction mode during several successive control steps.

More generally, the relevance of the LQG approach is bound to depend heavily on the adequation (or lack thereof) between the state-space model upon which this prediction is based and the physical reality. What is truly remarkable is that controllers based on simple AR(1) models with diagonal transition matrix perform generally better than the standard integral action control. However, as noted in Section 5, more refined disturbance models are needed for higher performance.

Identification of stochastic disturbance models for AO control has become a critical issue. A first logical step is to investigate AR models of higher orders. For exoplanet detection, AR(2) tip-tilt models have been chosen for example for the instruments SPHERE (on the VLT) and GPI (Gemini planet imager), both being on 8-meter class telescopes [19**? ?**]. Prior physical knowledge on perturbation can be combined with input/output data from the AO loop itself, as in [**? ? ?**]. Models identified using subspace identification and solely from WFS measurement data have been proposed too [25**? ?**]. Subspace methods have also been used to identify a dynamical discrete-time model of the DM+WFS subsystem [**?**].

One could imagine more complex models combining temporal priors with data from the loop itself and measurements from external sources, e.g. telescope-mounted atmospheric radar/lidar devices. Such a combination of information of different nature will inevitably require more sophisticated tuning methods. This is a difficult and largely open problem. On the other hand, an attractive potential benefit would be to refine the perturbation models while reducing the risk of parameter overfit [36], and thereby to increase performance robustness.

Another related difficulty is that while it is reasonable to regard atmospheric turbulence as a stationary process with respect to the AO loop's own time scale, its correlation structure nevertheless evolves with atmospheric conditions. Because astronomical imaging involves very long exposure times, operational LQG AO controllers should likely be equipped with some kind of self-tuning mechanism, based on periodic partial re-identification of the disturbance model. When based on physical priors, these models can benefit from methods that estimate spatial and geometric parameters from WFS measurements, see e.g. [**?**].

A closely related issue is of course the robustness in performance of the proposed LQG design with respect to modeling uncertainty/variations. While this is clearly a difficult problem in a general reconstructed feedback framework, an important simplification occurs when the DM's dynamics are neglected, since the LQG solution only involves one Riccati equation (for the Kalman filter). Some preliminary performance sensitivity tests have been simulated in good seeing configurations on a classical AO system with AR(1) models, showing a good robustness to model errors [6]. In addition, such crude AR(1) models have been used successfully in experimental conditions where a Taylor type turbulence (a rotating phase screen) does not correspond to any AR(1) model. Preliminary tests in MCAO have shown good robustness to parameter variation, where matrices M_α^{tur}, M_β^{tur}, M_α^{cor} and M_β^{cor} may also be slightly different from the real geometric configuration of the AO system to be controlled. More general results of practical interest could however possibly be obtained through theoretical studies of performance robustness.

This also suggests that the assumption of negligible DM's response time should not be discarded too frivolously: while an optimal LQG controller based on a more refined DM's model should in theory yield optimal performance, it is also bound to result in a more complex and less robust design, raising possibly tricky implementation issues – particularly if no real-time measurements of the DM's internal variables are available. Neglecting DM's dynamics may thus remain a sensible engineering trade-off for a number of AO systems.

This does not mean that the theoretical and practical issues raised by DM's dynamics should not be actively pursued. Indeed, new generation instruments include DMs with diameters in the order of one meter or more, which is about ten times the size of current DMs. These future DMs are generally expected to run at frequencies around 1 kHz. As an illustration, the future secondary mirror on the VLT will have 1100 actuators distributed over a 1.12 meter diameter surface, with a minimum sampling frequency of 500 Hz, to be increased to about 1 kHz. In the European Extremely Large Telescope project, a DM with a diameter of 2.5 m and between 5000 and 10000 piezoelectric actuators is envisioned, together with smaller DMs between 40 and 60 cm. The American Thirty Meter Telescope project features DMs in the order of 30 to 80 cm, with 61×61 to 121×121 actuators.

XXX Local control, model ident using FEM (ref GEPI), pb des modes propres.

Thus, one can state fairly confidently that at some point in the future, DM's dynamics, including non-linearities, will matter, if only because the long-term trend towards ever bigger DMs operating at ever higher sampling rates, and possibly made mechanically more flexible by the

[6] 50% variation of Σ_ϕ leads to less than 0.2 % SR loss, 50% variation of V leads to negligible loss, 50% variation of Σ_w leads to less than 0.1 % SR loss

use of innovative actuator technologies, has set the AO loops' disturbance rejection bandwidth and the set of DM's resonance frequencies on an implacable collision course.

Acknowledgements The authors would like to thank Carlos Correia da Silva for his useful comments.

A Inter-sampling variance

The inter-sampling variance $\varepsilon^2_{\text{sampl}}$ is un incompressible error term due to the application of the control through a zero-order hold. It depends only on the turbulence characteristics and sampling frequency of the AO loop. Take its definition in (6),

$$\varepsilon^2_{\text{sampl}} = \lim_{K\to+\infty} \frac{1}{K} \sum_{k=0}^{K-1} \left(\frac{1}{T} \int_{kT}^{(k+1)T} \left\| \phi^{\text{tur}}(t) - \phi^{\text{tur}}_{k+1} \right\|^2 \mathrm{d}t \right), \tag{A.1}$$

with the following hypothesis: the continuous variable $\phi^{\text{tur}}(t)$ is a weak stationary stochastic process, with known correlation matrix $\mathrm{E}(\phi^{\text{tur}}(t)\phi^{\text{tur}}(t+\tau)^{\text{t}}) = \Sigma_{\phi^{\text{tur}}}(\tau)$. Let us define $\psi^{\text{tur}}(t)$ as the filtered of $\phi^{\text{tur}}(t)$ by the filter with impulse response

$$h(t) = \frac{1}{T} \text{ for } 0 \le t < T. \tag{A.2}$$

The filtered process $\psi^{\text{tur}}(t)$ is thus obtained as

$$\psi^{\text{tur}}(t) = \frac{1}{T} \int_0^T \phi^{\text{tur}}(t-u)\mathrm{d}u. \tag{A.3}$$

The discrete process $\phi^{\text{tur}}_k = \frac{1}{T}\int_{(k-1)T}^{kT} \phi^{\text{tur}}(t)\mathrm{d}t$ is obtained directly as the sampling of $\psi^{\text{tur}}(t)$ at instants kT, that is $\phi^{\text{tur}}_k = \psi^{\text{tur}}(kT)$.

Develop (A.1) and invoke stationarity for the last line:

$$\begin{aligned} \varepsilon^2_{\text{sampl}} &= \mathrm{E}\left(\|\phi^{\text{tur}}(t)\|^2\right) - \mathrm{E}\left(\|\psi^{\text{tur}}(kT)\|^2\right), \\ &= \mathrm{E}\left(\|\phi^{\text{tur}}(t)\|^2\right) - \mathrm{E}\left(\|\psi^{\text{tur}}(t)\|^2\right). \end{aligned} \tag{A.4}$$

We thus have $\varepsilon^2_{\text{sampl}} = \text{trace}\Sigma_{\phi^{\text{tur}}}(0) - \text{trace}\Sigma_{\psi^{\text{tur}}}(0)$, where the correlation matrix $\Sigma_\psi(\tau) = \mathrm{E}(\psi^{\text{tur}}(t)\psi^{\text{tur}}(t+\tau)^{\text{t}})$ is available as $\Sigma_{\phi^{\text{tur}}}$ is known. By Parseval's theorem,

$$\varepsilon^2_{\text{sampl}} = \text{trace}\int_{-\infty}^{+\infty} S_{\phi^{\text{tur}}}(f)\mathrm{d}f - \text{trace}\int_{-\infty}^{+\infty} S_{\psi^{\text{tur}}}(f)\mathrm{d}f, \tag{A.5}$$

with $S_{\phi^{\text{tur}}}$ and $S_{\psi^{\text{tur}}}$ the PSDs of $\phi^{\text{tur}}(t)$ and $\psi^{\text{tur}}(t)$ respectively. From (A.2) and (A.3), $S_{\psi^{\text{tur}}}$ is computed as

$$S_{\psi^{\text{tur}}}(f) = |H(f)|^2 S_{\phi^{\text{tur}}}(f) = |\text{sinc}(\pi f T)|^2 S_{\phi^{\text{tur}}}(f), \tag{A.6}$$

so that $\varepsilon^2_{\text{sampl}} = \text{trace}\int_{-\infty}^{+\infty}(I - |\text{sinc}(\pi f T)|^2)S_{\phi^{\text{tur}}}(f)\mathrm{d}f$. For a phase expressed on a Zernike basis, the inter-sampling variance takes the form

$$\varepsilon^2_{\text{sampl}} = \sum_{i=1}^{\infty} \int_{-\infty}^{+\infty} (1 - |\text{sinc}(\pi f T)|^2) S_{a_i}(f))\mathrm{d}f, \tag{A.7}$$

where the S_{a_i}'s are the PSDs of the Zernike coefficients of the turbulent phase, which can be computed using standard priors [9]. The evaluation of (A.7) with respect to different atmospheric conditions can be found in [46].

References

[1] Horace W. Babcock. The possibility of compensating astronomical seeing. *Pub. Astron. Soc. Pacific*, 65:229, 1953.

[2] Lucie Baudouin, Christophe Prieur, Fabien Guignard, and Denis Arzelier. Robust control of a bimorph mirror for adaptive optics system. *J. Applied Optics*, 47(20):3637–3645, 2008.

[3] Jacques M. Beckers. Increasing the size of the isoplanatic patch with multiconjugate adaptive optics. In M.-H. Ulrich, editor, *Very Large Telescopes and their Instrumentation*, volume 2 of *ESO Conference and Workshop Proceedings*, pages 693–703, Garching Germany, March 1988. ESO.

[4] Jacques M. Beckers. Multi-conjugate adaptive optics: experiments in atmospheric tomography. In P. Wizinowich, editor, *Adaptive Optical Systems Technology*, volume 4007, page 1056, Bellingham, Washington, 2000. Proc. Soc. Photo-Opt. Instrum. Eng., SPIE.

[5] Dennis S. Bernstein, LLawrence D. Davis, and Scott W. Greeley. The optimal projection equations for fixed-order sampled-data dynamic compensation with computational delay. *IEEE Trans. Autom. Control*, AC-31:859–862, 1986.

[6] Tongwen Chen and Bruce A. Francis. *Optimal sampled-data control systems.* Springer-Verlag, London, 1995.

[7] Alessandro Chiuso, Riccardo Muradore, and Enrico Marchetti. Dynamic calibration of adaptive optics systems: A system identification approach. pages 750–755, Cancun, 2008. 47th IEEE Conference on Decision and Control.

[8] Jean-Marc Conan and Gérard Rousset, editors. *Multi-Conjugate Adaptive Optics for very large telescopes*, volume 6(10) of *Comptes Rendus Physique*, pages 1035–1194. Académie des Sciences / Elsevier Masson, France, dec 2005.

[9] Jean-Marc Conan, Gérard Rousset, and Pierre-Yves Madec. Wave-front temporal spectra in high-resolution imaging through turbulence. *J. Opt. Soc. Am. A*, 12(12):1559–1570, 1995.

[10] Carlos Correia, Henri-François Raynaud, Caroline Kulcsár, and Jean-Marc Conan. Globally optimal minimum-variance control in adaptive optical

systems with mirror dynamics. In Claire E. Max Norbert Hubin and Peter L. Wizinowich, editors, *Adaptive Optics Systems: Real time Control and Algorithms*, volume 7015, page 70151F. Proc. Soc. Photo-Opt. Instrum. Eng., 2008.

[11] Chris E. Coulman, Jean Vernin, and Alain Fuchs. Optical seeing-mechanism of formation of thin turbulent laminae in the atmosphere. *Appl. Opt.*, 34:5461–5474, 1995.

[12] Carlos E. De Souza and Graham C. Goodwin. Intersample variances in discrete minimum variance control. *IEEE Trans. Autom. Control*, AC-29:759–761, 1984.

[13] Caroline Dessenne, Pierre-Yves Madec, and Gérard Rousset. Optimization of a predictive controller for closed-loop adaptive optics. *Appl. Opt.*, 37(21):4623–4633, jul 1998.

[14] Robert H. Dicke. Phase-contrast detection of telescope seeing and their correction. *Astron. J.*, 198(4):605–615, 1975.

[15] Brent Ellerbroek and Curtis R. Vogel. Simulations of closed-loop wavefront reconstruction for multi-conjugate adaptive optics on giant telescopes. In R. K. Tyson and M. Lloyd-Hart, editors, *Astronomical Adaptive Optics Systems and Applications*, volume 5169, pages 206–217. Proc. Soc. Photo-Opt. Instrum. Eng., SPIE, 2003.

[16] Thierry Fusco. *Correction partielle et anisoplanétisme en Optique Adaptative : traitements* a posteriori *et Optique Adaptative Multiconjuguée.* PhD thesis, Université de Nice Sophia-Antipolis, 2000.

[17] Thierry Fusco, Cyril Petit, Gérard Rousset, Jean-Marc Conan, and Jean-Luc Beuzit. Closed-loop experimental validation of the spatially filtered shack-hartmann concept. *J. Opt. Soc. Am. A*, 30(11):1255–1257, 2005.

[18] Thierry Fusco, Gérard Rousset, Jean-Luc Beuzit, David Mouillet, Kjetil Dohlen, Rodolph Conan, Cyril Petit, and Guillaume Montagnier. Conceptual design of an extreme ao dedicated to extrasolar planet detection by the vlt-planet finder intrument. In *Astronomical Adaptive Optics Systems and Applications II*, volume 5903, pages 59030K.1–59030K.12. Proc. Soc. Photo-Opt. Instrum. Eng., SPIE, 2005. 2005, San Diego, USA.

[19] Thierry Fusco, Gérard Rousset, Jean-François Sauvage, Cyril Petit, Jean-Luc Beuzit, Kjetil Dohlen, Julien Charton, Magali Nicolle, Markus Kasper, Pierre Baudoz, and Pascal Puget. High-order adaptive optics requirements for direct detection of extrasolar planets: Application to the sphere instrument. *oe*, 14:7515–7534, 2006.

[20] D. Gavel and D. Wiberg. Toward strehl-optimizing adaptive optics controllers. In P. L. Wizinowich, editor, *Adaptive Optical Systems Technologies II*, volume 4839, pages 890–901. Proc. Soc. Photo-Opt. Instrum. Eng., SPIE, 2002.

[21] Éric Gendron and Pierre Lena. Astronomical Adaptive optics I. modal control optimization. *Astronomy and Astrophysics*, 291:337–347, 1994.

[22] Éric Gendron and Pierre Lena. Astronomical Adaptive Optics II. experimental results of an optimize modal control. *Astronomy and Astrophysics*, 111:153–167, 1994.

[23] James S. Gibson and Brent L. Ellerbroek. Adaptive optics wave-front correction by use of adaptive filtering and control. *Applied Optics*, 39(16):2525–2538, 2000.

[24] François Hammer, Mathieu Puech, François Assémat, Éric Gendron, Frédéric Sayede, Philippe Laporte, Michel Marteaud, Arnaud Liotard, and Frédéric Zamkotsian. Falcon: a concept to extend adaptive optics corrections to cosmological fields. In A. L. Ardeberg and T. Andersen, editors, *Second Backaskog Workshop on Extremely Large Telescopes*, volume 5382, pages 727–736. Proc. Soc. Photo-Opt. Instrum. Eng., 2004.

[25] Karel Hinnen, Michel Verhaegen, and Niek Doelman. Exploiting the spatiotemporal correlation in adaptive optics using data-driven h2-optimal control. *J. Opt. Soc. Am. A*, 24:1714–1725, 2007.

[26] Karel Hinnen, Michel Verhaegen, and Niek Doelman. A data-driven h_2-optimal control approach for adaptive optics. *IEEE Trans. on Control Systems Technology*, 16(3):381–395, 2008.

[27] Eric W. Justh, Perinkulam Sambamurthy Krishnaprasad, and Mikhail A. Vorontsov. Analysis of a high-resolution optical wave-front control system. *Automatica*, 40(7):1129–1141, July 2004.

[28] A. N. Kolmogorov. Local structure of turbulence in incompressible fluids with very high reynolds number. *Dokl. Akad. Nauk. SSSR*, 30(4):301–305, 1941.

[29] Caroline Kulcsár, Henri-François Raynaud, Cyril Petit, and Jean-Marc Conan. Minimum variance control in presence of actuator saturation in adaptive optics. In Claire E. Max Norbert Hubin and Peter L. Wizinowich, editors, *Adaptive Optics Systems: Real time Control and Algorithms*, volume 7015, page 70151G. Proc. Soc. Photo-Opt. Instrum. Eng., 2008.

[30] Caroline Kulcsár, Henri-François Raynaud, Cyril Petit, Jean-Marc Conan, and Patrick Viaris de Lesegno. Optimal control, observers and integrators in adaptive optics. *Optics Express*, 14(17):7463–8012, 2006.

[31] Vladim`Kučera. *Analysis and design of discrete linear control systems.* Series in Systems and Control Engineering. Prentice Hall, Herfordshire, UK, 1991.

[32] Huibert Kwakernaak and Raphael Sivan. *Linear Optimal Control Systems.* Wiley-Interscience. John Wiley & Sons, Inc., Wiley-Interscience, New-York, 1972.

[33] Pierre Le Gall, Christophe Prieur, and Lionel Rosier. On the control of a bimorph mirror. Cachan, France, 2006. IFAC Workshop on Control Applica-

tions of Optimization.

[34] Brice Le Roux, Jean-Marc Conan, Caroline Kulcsár, Henri-François Raynaud, Laurent M. Mugnier, and Thierry Fusco. Optimal control law for multiconjugate adaptive optics. In Peter L. Wizinowich and Domenico Bonaccini, editors, *Adaptive Optical System Technology II*, volume 4839, pages 878–889, Hawaii, USA, 2002. Proc. Soc. Photo-Opt. Instrum. Eng., SPIE.

[35] Brice Le Roux, Jean-Marc Conan, Caroline Kulcsár, Henri-François Raynaud, Laurent M. Mugnier, and Thierry Fusco. Optimal control law for classical and multiconjugate adaptive optics. *J. Opt. Soc. Am. A*, 21(7):1261–1276, 2004.

[36] Lennart Ljung. *System Identification, Theory for the User*. Prentice Hall, Upper Saddle River, New Jersey, 2nd edition edition, 1999.

[37] Douglas P. Looze. Minimum variance control structure for adaptive optics systems. *J. Opt. Soc. Am. A*, 23(3):603–612, March 2006.

[38] Douglas P. Looze. Discrete-time model of an adaptive optics system. *J. Opt. Soc. Am. A*, 24(9):2850–2863, September 2007.

[39] Douglas P. Looze. Linear-quadratic-gaussian control for adaptive optics systems using a hybrid model. *J. Opt. Soc. Am. A*, 26(1):1–9, January 2009.

[40] Douglas P. Looze, Orhan Beker, Markus Kasper, and Stephan Hippler. Optimal compensation and implementation for adaptive optics systems. volume 2, pages 1715–1720, Phoenix, AZ, USA, 1999. IEEE Conf. on Decision and Control.

[41] Douglas P. Looze, Markus Kasper, Stefan Hippler, Orhan Beker, and Robert Weiss. Optimal compensation and implementation for adaptive optics systems. *Experimental Astronomy*, 15:67–88, 2003.

[42] David W Miller and Simon C. O. Grocott. Robust control of the multiple mirror telescope adaptive secondary mirror. *Opt. Eng.*, 38(8):1276–1287, 1999.

[43] Robert J. Noll. Zernike polynomials and atmospheric turbulence. *J. Opt. Soc. Am.*, 66(3):207–211, 1976.

[44] Randall N. Paschall and David J. Anderson. Linear Quadratic Gaussian control of a deformable mirror adaptive optics system with time-delayed measurements. *Applied Optics*, 32(31):6347–6358, 1993.

[45] Randall N. Paschall, Mark A. Von Bokern, and B. M Welsh. Design of a linear quadratic gaussian controller for an adaptive optics system. volume 2, pages 1761–1769, Brighton, UK, December 1991. IEEE Conf. on Decision and Control.

[46] Cyril Petit. *Étude de la commande optimale en OA et OAMC, validation numérique et expérimentale*. PhD thesis, Université Paris 13, 2006.

[47] Cyril Petit, Jean-Marc Conan, Caroline Kulcsár, and Henri-François Raynaud. Linear quadratic gaussian control for adaptive optics and multiconjugate adaptive optics: experimental and numerical analysis. *J. Opt. Soc. Am. A*, 26(6):1307–1325, 2009.

[48] Cyril Petit, Jean-Marc Conan, Caroline Kulcsár, Henri-François Raynaud, and Thierry Fusco. First laboratory validation of vibration filtering with lqg control law for adaptive optics. *Optics Express*, 16(1):87–97, 2008.

[49] Cyril Petit, Jean-Marc Conan, Caroline Kulcsár, Henri-François Raynaud, Thierry Fusco, Joseph Montri, and Didier Rabaud. Optimal control for multi-conjugate adaptive optics. *Comptes Rendus de l'Académie des Sciences, Physique*, 6:1059–1069, 2005.

[50] Piotr Piatrou and Michael C. Roggemann. Performance study of kalman filter controller for multiconjugate adaptive optics. *Appl. Opt.*, 46(9):1446–1455, March 2007.

[51] Lisa A. Poyneer and Bruce A. Macintosh. Spatially filtered wave-front sensor for high-order adaptive optics. *J. Opt. Soc. Am. A*, 21:810–819, 2004.

[52] Lisa A. Poyneer, Bruce A. Macintosh, and Jean-Pierre Véran. Fourier transform wavefront control with adaptive prediction of the atmosphere. *J. Opt. Soc. Am. A*, 24:2645–2660, 2007.

[53] Lisa A. Poyneer and Jean-Pierre Véran. Predictive wavefront control for adaptive optics with arbitrary control loop delays. *J. Opt. Soc. Am. A*, 25:1486–1496, 2008.

[54] Fernando Quirós-Pacheco. *Reconstruction and Control Laws for Multi-conjugate Adaptive Optics in Astronomy*. PhD thesis, Imperial College, London, 2006.

[55] Henri-François Raynaud, Caroline Kulcsár, Carlos Correia da Silva, and Jean-Marc Conan. Multirate lqg ao control. In *Adaptive Optics Systems*, volume 7015, page 701538, Marseille, 2008. Proc. Soc. Photo-Opt. Instrum. Eng., SPIE.

[56] François Roddier. *Adaptive Optics in Astronomy*. Cambridge, Cambridge University Press, 1999.

[57] Gérard Rousset, Jean-Claude Fontanella, Pierre Y. Kern, Pierre Gigan, François Rigaut, Pierre L, Corinne Boyer, Pascal Jagourel, Jean-Paul Gaffard, and Fritz Merkle. First diffraction-limited astronomical images with adaptive optics. *Astron. Astrophys.*, 230:29–32, 1990.

[58] Don Wiberg and Donald T. Gavel. A spatial nondynamic lqg controller: Part i, application to adaptive optics. volume 3, pages 3326–3332, Nassau, Bahamas, Dec. 2004. IEEE Conf. on Decision and Control.

[59] Allan Wirth, Joseph Navetta, Douglas Looze, Stefan Hippler, Andreas Glindemann, and Donald Hamilton. Real-time modal control implementation for adaptive optics. *Appl. Opt.*, 37(21):4586–4597, jul 1998.

[60] Kemin Zhou, Keith Glover, and John Doyle. *Robust and Optimal Control*. Prentice Hall, Upper Saddle River, 1996.

[61] Guchuang Zhu, Jean Lévine, Laurent Praly, and Yves-Alain Peter. Flatness-based control of electro-

statically actuated mems with application to adaptive optics: a simulation study. *Journal of Microelectromechanical Systems*, 1(5):1165–1174, 2006.