A correlation-locking adaptive filtering technique for minimum variance integral control in adaptive optics

V. Deo^{1,2}, É. Gendron¹, F. Vidal¹, M. Rozel¹, A. Sevin¹, F. Ferreira¹, D. Gratadour^{1,3}, N. Galland¹, and G. Rousset¹

¹ LESIA, Observatoire de Paris, Univ. PSL, CNRS, Sorbonne Univ., Univ. de Paris, 5 pl. Jules Janssen, 92195 Meudon, France
e-mail: vincent.deo@obspm.fr

² National Astronomical Observatory of Japan, Subaru Telescope, 650 North A'ohōkū Place, Hilo, HI 96720, USA

³ Research School of Astronomy and Astrophysics, Australian National University, Canberra, ACT 2611, Australia

Received 22 December 2020 / Accepted 17 March 2021

ABSTRACT

We propose the correlation-locking optimization scheme (CLOSE), a real-time adaptive filtering technique for adaptive optics (AO) systems controlled with integrators. CLOSE leverages the temporal autocorrelation of modal signals in the controller telemetry and drives the gains of the integral command law in a closed servo-loop. This supervisory loop is configured using only a few scalar parameters and automatically controls the modal gains to closely match transfer functions achieving minimum variance control. This optimization is proven to work throughout the range of noise and seeing conditions relevant to the AO system. This technique was designed while preparing the high-order AO systems for extremely large telescopes, in particular to tackle the optical gain (OG) phenomenon. This is a sensitivity reduction induced by on-sky residuals and is a prominent issue with pyramid wavefront sensors (PWFS). CLOSE follows upon the linear modal compensation approach to OG that was previously demonstrated to substantially improve AO correction with high-order PWFS systems. Operating on modal gains through multiplicative increments, CLOSE naturally compensates for the recurring issue of unaccounted sensitivity factors throughout the AO loop. We present end-to-end simulations of the MICADO instrument single-conjugate AO to demonstrate the performances and capabilities of CLOSE. We demonstrate that a single configuration provides an efficient and versatile optimization of the modal integrator while accounting for OG compensation and while providing significant robustness to transient effects impacting the PWFS sensitivity.

Key words. instrumentation: adaptive optics – techniques: high angular resolution – telescopes

1. Introduction

Thanks to its extreme sensitivity as compared to other general-purpose wavefront sensors (WFSs) for astronomical adaptive optics (AO), the pyramid WFS (PWFS; [Ragazzoni 1996](#)) has been the design choice included in all first-light AO systems for the three upcoming extremely large telescopes (ELTs; [Tamai et al. 2018](#); [Fanson et al. 2018](#); [Liu & Sanders 2018](#)) as well as for a number of high-performance systems on 8–10 m telescopes ([Esposito et al. 2010](#); [Guyon et al. 2010](#); [Schatz et al. 2018](#)).

In recent years, the community has shown interest in tackling the nonlinearity of the PWFS, the so-called optical gain (OG), which is an on-sky sensitivity reduction induced by the limited dynamic range of the sensor and can be modeled as a function of wavefront spatial frequency, with a magnitude depending on residual wavefront conditions, and therefore on ongoing turbulence statistics ([Costa 2005](#); [Korkiakoski et al. 2008a](#); [Deo et al. 2018a](#)). When unmeasured and not compensated for, OG is a significant hindrance that prevents the application of many algorithms relying upon the linearity of the servo-loop. Two extremely common but critical examples are (1) the explicit estimation of the temporal transfer function, which is necessary for applying modal gain optimization techniques ([Gendron & Léna 1994](#); [Dessenne et al. 1998](#)), and (2) the proper subtraction of the noncommon path aberration setpoint (examined in, e.g., [Esposito et al. 2020](#)). For setpoint subtraction, the PWFS triggers a divergent positive feedback when attempting convergence

to a wavefront setpoint whose gradient exceeds the modulation radius.

Our previously proposed modal OG compensation pipeline ([Deo et al. 2019a](#)) demonstrated significant performance improvements for different seeing conditions, but had some limitations: It required introducing probe signals on the deformable mirror (DM), it required a significant amount of preliminary modeling and computations, and it did not incorporate any modal variance minimization technique depending on noise level variations. We propose in the present paper a novel algorithm, a self-regulating and entirely automated real-time modal gain controller, that resolves these three shortcomings: the correlation-locking optimization scheme (CLOSE). The method is inspired by [Montera et al. \(2018\)](#), proposing to apply neural networks in particular for tip-tilt sensitivity tracking and compensation. Preliminary results with CLOSE were published in [Deo et al. \(2019b\)](#).

CLOSE monitors the integrator overshoot through the temporal autocorrelation of modal WFS measurements and drives modal gains through real-time multiplicative updates. This allows tracking and optimizing the sole temporal properties of the loop integrator. The OG compensation multipliers are automatically factored in the gain that is set into the command law and do not require separate explicit computations. The resulting steady-state command law may then be optimized regardless of the OG sensitivity reduction and can easily be tuned so as to compare to the minimum variance (MV) control of [Gendron & Léna \(1994\)](#).

Fig. 1. General modeling of the SCAO loop showing the WFS, real-time computer, and DM. The WFS is modeled using the confusion matrix framework (Sect. 2.1). The confusion matrix $\mathbf{A}$ is a random variable depending upon the residual wavefront ϕ_{Res} .

While we started designing CLOSE in the context of the single-conjugate AO (SCAO) module of the MICADO instrument (Davies et al. 2018; Clénet et al. 2019), the technique naturally extends to a range of AO systems, namely for nonlinear sensors performing well under a locally valid, modal, linearized description, such as the PWFS. This includes other Fourier-like wavefront sensors (Fauvarque et al. 2019; Chambouleyron et al. 2020) as well as quad-cell Shack–Hartmanns or the adaptation of a diffraction-limited, point-source AO calibration into an adequate command law for extended natural sources or elongated laser stars, etc.

This paper is organized as follows: in Sect. 2 we present our formal model for analyzing the AO loop dynamics when using a WFS prone to optical gain. In Sect. 3 we present the rationale for the CLOSE servo-loop as well as semi-analytical demonstrations showing the achieved correspondence with minimum variance integrators. Then we present in Sect. 4 two proposed implementations, real-time and offline, and briefly lay out the computational requirements associated with the use of CLOSE. We show in Sect. 5 the results of end-to-end numerical simulations, showing convergence, improved performance, and the breadth of applicable conditions for CLOSE in the MICADO SCAO system. Finally, Sect. 6 includes some discussions of possible extensions and limitations of the scheme.

2. Modeling the AO loop

We show in Fig. 1 the schematic of a SCAO loop that is considered throughout this paper. This model that was designed to conveniently describe the OG effect on PWFSs operating in low Strehl regimes. We recall here the key elements relevant to this study. A more extensive description of this approach may be found in Deo et al. (2019a).

While the wavefronts defined in Fig. 1, ϕ_{Atm} , ϕ_{DM} , and ϕ_{Res} , are continuous functions over the telescope aperture, but are here implicitly meant as their decomposition over a control basis of the DM: $(\phi_1, \dots, \phi_N)$, where N is the number of controlled modes, plus an additional fitting component beyond the DM capabilities. The left half of Fig. 1 is represented in this modal space. The right half, where the measurement noise $\mathbf{p}$ is introduced, is in the WFS measurement space. For a Shack–Hartmann, this would be the centroid displacement space, or for a PWFS, the space of either the normalized pixels or the gradient-like slopes maps.

The WFS is represented by the OG-describing matrix $\mathbf{A}$ (Sect. 2.1) and the transformation from modal decomposition to measurement $\mathbf{dWFS}$: the modal interaction matrix of the wavefront sensor, defined as the Jacobian of the WFS response around the flat wavefront, that is, computed using infinitesimal push-

pulls. Wavefront reconstruction is performed by matrix-vector multiplication, with the matrix $\mathbf{Rec}$ computed as the generalized inverse of $\mathbf{dWFS}$, assuming the latter is adequately conditioned. Finally, Fig. 1 shows that the temporal control of the loop is operated through a modal integrator using a gain vector $\mathbf{G} = [G_i]_{1 \leq i \leq N}$.

2.1. Optical gain: The confusion matrix model

In the general case, the small-signal response of a nonlinear sensor is modified by the presence of atmospheric wavefront residuals. The WFS Jacobian $\mathbf{dWFS}(\phi_{\text{Res}})$ around any nonzero setpoint ϕ_{Res} may be written as $\mathbf{dWFS}(\phi_{\text{Res}}) = \mathbf{A} \cdot \mathbf{dWFS}$, as shown in Fig. 1. In this description, the modal space operator $\mathbf{A}$ describes a local mixing of model components around ϕ_{Res} as compared to the calibrated response, and we therefore call $\mathbf{A}$ the modal confusion matrix. The operator $\mathbf{A}$ generally is a random variable dependent on ϕ_{Res} . In the simplest case, for instance, a Shack–Hartmann WFS with uniform centroid gain, the matrix $\mathbf{A}$ is a scaled identity matrix.

For the PWFS case, it has been shown that the confusion matrix has some reasonable properties when described on an appropriate modal basis. These properties are the foundation of optical gain modal compensation for the PWFS (Korkiakoski et al. 2008b). In previous work (Deo et al. 2018a, 2019a), we performed a thorough numerical assessment of the fluctuations of $\mathbf{A}$ when the spatial power spectrum density (PSD) of ϕ_{Res} is stationary. These analyses were performed using a Karhunen–Loève (KL) basis orthonormalized on the DM (Ferreira et al. 2018a). This basis is made of modes containing an isotropic mix of spatial frequencies of a single norm, sorted by increasing frequency. The last ~ 300 modes (out of 4300 total modes for the ELT), whose structure is affected by the DM cut-off, contain a variety of waffles. We use this basis for all purposes in this paper.

When a decomposition is used on our DM KL basis, we have previously demonstrated that (1) $\mathbf{A}$ is essentially diagonal for low-order modes, which bear most of the power of the atmospheric turbulence; (2) that its diagonal coefficients vary by no more than a few percent for a given set of wavefronts ϕ_{Res} of identical PSD, a property in accordance with convolutional PWFS descriptions (Fauvarque et al. 2019; Chambouleyron et al. 2020); and (3) that the off-diagonal portion of $\mathbf{A}$, while non-negligible for high-order ϕ_i , is of negligible statistical average.

Using this basis, which statistically diagonalizes $\mathbf{A}$, enables the modal gain compensation strategy, as shown in the Reconstruction block of Fig. 1. The modal gains set in $\mathbf{G}$ cover two functionalities: First, they include the compensation of $\mathbf{A}$ for

Fig. 2. Diagonalization of the AO loop model (Fig. 1) for mode i under the approximation that $\mathbf{A}$ is diagonal and is thus reduced to its i th diagonal term α_i .

ongoing turbulent conditions; and second, they define the transfer function gain such that the modal integrators exhibit appropriate rejection levels. Both of these terms are ambiguously mixed into the coefficients of $\mathbf{G}$ set into the controller. This ambiguity affects the dynamical modeling of the system, which we describe in Sect. 2.3, as well as the optimization strategies explained in Sect. 3. Its relation to proper noncommon path aberration compensation, which requires independently identifying the compensation of $\mathbf{A}$, is discussed in Sect. 6.2. The stability of $\mathbf{A}$ against the PSD of ϕ_{Res} ensures that adequate G_i values vary with the same temporal scale as the descriptive statistical parameters (r_0 , L_0 , $C_n^2(h)$, ...) of the turbulence.

2.2. Diagonalizing the loop model

Following through with the approximation of statistical diagonality of $\mathbf{A}$, Fig. 1 can then be simplified to the flowchart shown in Fig. 2 because the reference interaction and reconstruction matrices $\mathbf{dWFS}$ and $\mathbf{Rec}$ reduce to an identity matrix. Figure 2 applies as one of N decoupled servo-loops for each of the controlled modes, with $\mathbf{A}$ reduced to its i th diagonal coefficient α_i . We call α_i the modal sensitivity reduction coefficient because $0 < \alpha_i \leq 1$ stands for a PWFS. With α_i varying only with the statistical properties of the turbulence, it can be considered as a constant for the purposes of analyzing the model in Fig. 2 for stationary or slowly evolving atmospheric conditions. With these hypotheses, Fig. 2 is no more than a most classical linear integral feedback loop, but for the unknown and unmeasured α_i .

Furthermore, our choice of basis is descriptively similar to atmospheric KL modes (Dai 1996) for the modes that bear the most of the turbulence power. Thus, the N modal loops as shown in Fig. 2 may be approximated as uncorrelated, with their independent optimizations resulting in a global one.

2.3. Transfer functions

Each of the N modal servo-loops is entirely described by a small number of parameters: the temporal spectra of the turbulence and noise for the i th mode, and the α_i (unknown) and G_i (known) scalars. The atmospheric temporal spectrum for a single KL mode is well described in the literature (Conan et al. 1995; Gendron 1995) for Kolmogorov or Von Kármán turbulence,

$$|\hat{\phi}_{\text{Atm}}^i(f)|^2 \propto f^0 \text{ for } f < f_i, \text{ and } \propto f^{-17/3} \text{ for } f > f_i, \quad (1)$$

where f denotes temporal frequencies, $\hat{\cdot}$ temporal Fourier transforms, and f_i is the mode cutoff frequency, given by the wind

speed and the mode radial order. We use the spectrum of Eq. (1) as a reference template for the analyses described in Sect. 3.

The minimization objective we consider for integral control optimization is the variance of the residual,

$$\text{Var}_k(\phi_{\text{Res}}^i[k]) = \int_f |\hat{\phi}_{\text{Res}}^i(f)|^2 df, \quad (2)$$

where k is the temporal sample index. The value of the $\phi_{\text{Res}}^i[k]$ cannot be accessed directly, however, but only those of the measurements $m_i[k]$. To perform semianalytical computations of how to achieve this minimization, as we do in Sect. 3.3, it is necessary to introduce the various transfer functions between the quantities involved.

These can easily be obtained (e.g., Madec 1999) from Fig. 2, and we here recall their expressions in a time-sampled framework. A discussion of continuous versus discrete time approaches to AO is proposed in Appendix A. We only note here that discrete-time is valid (Kulcsár et al. 2006) because (1) temporal aliasing is negligible, which is the case from the rapid decrease of Eq. (1); and (2) frequencies involved are well below Nyquist. We define the shorthand (Eq. (A.8)), where T is the sampling period and $j^2 = -1$,

$$h(f; g) = \left[1 + g \cdot \frac{\exp(-2j\pi f(T + \tau))}{1 - \exp(-2j\pi fT)} \right]^{-1}, \quad (3)$$

using which, we express the transfer functions, where the subscripts identify the input (either ϕ_{Atm}^i or the modal noise n_i) and output (ϕ_{Res}^i or the modal measurement m_i) considered,

$$h_{\text{Atm} \rightarrow \text{Res}}^i = h(f; \alpha_i G_i), \quad (4)$$

$$h_{n \rightarrow \text{Res}}^i = -\frac{1}{\alpha_i} \frac{\exp(-2j\pi f(T + \tau))}{1 - \exp(-2j\pi fT)} h(f; \alpha_i G_i), \quad (5)$$

and

$$h_{\text{Atm} \rightarrow m}^i = \alpha_i h(f; \alpha_i G_i) \quad (6)$$

$$h_{n \rightarrow m}^i = h(f; \alpha_i G_i). \quad (7)$$

These transfer functions are formally similar to those of an AO loop that is not affected by optical gain ($\alpha_i = 1$), except for an effective noise amplification, as seen from the $-1/\alpha_i$ in Eq. (5). It follows from Eqs. (6) and (7) that the Fourier transform of the measurements can be written as

$$\hat{m}_i(f) = h(f; \alpha_i G_i) \left(\alpha_i \hat{\phi}_{\text{Atm}}^i(f) + \hat{n}_i(f) \right). \quad (8)$$

Equation (8) emphasizes that for a given turbulence PSD, $\hat{m}_i(f)$ is entirely defined by two parameters: the transfer function gain $\alpha_i G_i$, and the sensitivity-adjusted signal-to-noise ratio (S/N), noted σ_i ,

$$\sigma_i^2 = \alpha_i^2 \frac{\text{Var}(\hat{\phi}_{\text{Atm}}^i)}{\text{Var}(\hat{n}_i)}. \quad (9)$$

Minimum variance control involves setting G_i such that $\alpha_i G_i$ is the minimizer of the quantity expressed in Eq. (2) for the given adjusted S/N. This either requires measuring α_i explicitly, as is done for instance in other optical gain compensation methods (Korkiakoski et al. 2008b; Esposito et al. 2015, 2020), or finding a proxy to indirectly infer the hidden value $\alpha_i G_i$, and adjust G_i adequately.

3. Correlation-locking scheme

When the modal sensitivity reductions α_i are unknown, the closed-loop measurements m_i still contain sufficient information to retrieve $h(f; \alpha_i G_i)$ and control G_i toward optimal rejection. We intend to perform this control without computing an extensive spectrum of the measurements and without performing an explicit estimation of the system response, but leveraging the short-term temporal autocorrelation of the measurements as a proxy for the effective modal loop gain $\alpha_i G_i$. We present in this section the correlation-locking optimization scheme (CLOSE), which drives the modal gains G_i in real-time toward a favorable solution for rejection. This section covers the foundation and steady-state solutions for the loop transfer functions achieved using CLOSE. We then present in Sect. 4 the implementation of CLOSE as a second-layer supervisory loop, taking the modal WFS measurements as inputs and operating on the gain vector.

3.1. Rationale: Using the loop resonance

For each mode, the loop described on Fig. 2 is a classical feedback loop with an integral controller of gain $\alpha_i G_i$ and total delay¹ $\tau + T$. The transfer function $h(f; g)$, as developed in Sect. 2.3, is a well-characterized high-pass filter, with a f^{+2} square modulus up to the roll-off, followed by a resonance peak located beyond. Some examples of $h(f; g)$, using $\tau = 2T$, are shown in Fig. 3 (top). $h(f; g)$ is a stable filter up to a maximum gain value $g = g_{\text{crit}}$, which depends only on the normalized latency $\frac{\tau}{T}$. As the gain g increases toward g_{crit} , the resonant peak sharpens and its amplitude increases. The central peak frequency increases and converges toward the critical frequency f_{crit} . In other terms, the lowest frequency pole of the Laplace transform associated with $h(f; g)$ is displaced toward the imaginary axis and eventually intercepts it at $s = 2j\pi f_{\text{crit}}$ for $g = g_{\text{crit}}$. Values of f_{crit} and g_{crit} with latency are provided in Table 1 for reference. The given general formula is demonstrated in Appendix A.

The behavior of the rejection peak as the transfer function approaches divergence is illustrated in Fig. 3, showing the PSDs for the atmospheric mode and the noise, and some examples of transfer functions and output spectra $|\hat{m}_i(f)|^2$ for three values of $g = \alpha_i G_i$. The latency is $\tau = 2T$, and the sensitivity adjusted S/N (Eq. (9)) was taken as $\sigma_i = 10$. Without loss of generality, we normalized the atmospheric spectrum at $|\alpha_i \hat{\phi}_{\text{Atm}}^i(f=0)|^2 = 1$.

The atmospheric spectrum, filtered by the high-pass transfer function, results in the leftmost peak of $|\hat{m}_i(f)|^2$ seen on Fig. 3 (bottom), at the modal cutoff frequency $f_i = 1$ Hz. Near 10 Hz, we see a transition from rejected turbulence, showing a $f^{-11/3}$ spectrum, into a regime of f^{+2} , as the input $\alpha_i \hat{\phi}_{\text{Atm}}^i + n_i$ is dominated by the measurement noise. This regime is followed by the rejection peak, which amplifies and shifts toward $f_{\text{crit}} = 50$ Hz as the gain increases.

3.2. Correlation locking: Steady-state objective

We assume that the AO latency τ has been calibrated and is a fixed known parameter; consequently, so is f_{crit} . For reasonable τ values, f_{crit} lies in the noise floor of the spectrum, well beyond the turbulence cutoff frequency. In its simplest form, our philosophy is to note that the amplitude of the resonant peak may be used as a proxy for the effective modal gain $\alpha_i G_i$.

Fig. 3. Typical spectral components for our AO semianalytical computations. *Top*: input spectra for turbulence and noise, and the measurement power transfer function $|h^i(f; \alpha_i G_i)|^2$ for three $\alpha_i G_i$ values. *Bottom*: corresponding measurement power spectra $|\hat{m}_i(f)|^2$. For this example, $\tau = 2T$, and $f_s = 500$ Hz, yielding $g_{\text{crit}} \approx 0.61$ and $f_{\text{crit}} = 50$ Hz.

Table 1. Parameters related to transfer function divergence depending on the latency of the system.

τ/T (frames)	f_{crit}/f_s	g_{crit}	δk_{crit} (frames)
0	1/2	2.0	1
1	1/6	1.0	3
2	1/10	≈ 0.618	5
$\frac{\tau}{T}$	$1/\left(4\frac{\tau}{T} + 2\right)$	$2 \sin\left(\pi \frac{f_{\text{crit}}}{f_s}\right)$	$2\frac{\tau}{T} + 1$

Notes. f_s is the sampling rate. $\delta k_{\text{crit}} = f_s/2f_{\text{crit}}$ is discussed in Sect. 3.2.

We deemed it particularly inconvenient, however, to attempt real-time estimations of the peak amplitude or structure, which calls for extensive buffer acquisitions, explicit PSD estimations, etc. The latter remains well within technical reach and may be done at later stages of this research.

Instead, we propose to use the autocorrelation (AC) of the modal measurements, noted $m_i^*[\delta k]$, because it provides another indirect measure of $\alpha_i G_i$. For reference, the AC curves $m_i^*[\delta k]$ at small time-shifts δk corresponding to the measurement spectra shown in the bottom panel of Fig. 3 are shown in Fig. 4. As $\alpha_i G_i$ increases and as the loop response approaches divergence, an oscillation of half-period nearing

$$\delta k_{\text{crit}} = \frac{f_s}{2f_{\text{crit}}} = 2\frac{\tau}{T} + 1 \text{ frames} \quad (10)$$

is superimposed on the typically wider bell-shaped AC curve. This oscillation is the correspondence in the AC domain of the resonance peak, growing in amplitude and converging to f_{crit} .

¹ Expected value of the delay between the occurrence of a perturbation $\phi_{\text{Atm}}(t)$ and the mean time of its correction on the DM.

Fig. 4. Autocorrelation functions $m_i^*[\delta k]$ corresponding to the measurement PSDs $|\hat{m}_i|^2$ shown in Fig. 3. $\delta k_{\text{crit}} = 5$ frames for $\frac{\tau}{T} = 2$.

In particular, the first minimum of the AC curve for a time-shift of δk_{crit} is a strongly marked signal that allows measuring the amount of resonance in the loop with a minimal latency of δk_{crit} frames. This value is the one metric that is monitored by CLOSE to implicitly register the value of the loop gain $\alpha_i G_i$,

$$m_i^*[\delta k_{\text{crit}}] = \frac{\sum_k m_i[k] m_i[k - \delta k_{\text{crit}}]}{\sum_k m_i[k]^2}, \quad (11)$$

which, as we show in Sect. 4, is also convenient to track with simple real-time estimators. Decreasing $\alpha_i G_i$ reduces the amplitude of the oscillation over the AC function, which tends to increase the value of $m_i^*[\delta k_{\text{crit}}]$, and the opposite occurs when $\alpha_i G_i$ is increased (within the stability limits).

Because the relation between $\alpha_i G_i$ and $m_i^*[\delta k_{\text{crit}}]$ is monotonic, we can act on G_i in order to lock the δk_{crit} correlation value onto a steady-state solution,

$$m_i^*[\delta k_{\text{crit}}] = r \in [-1, 1], \quad (12)$$

where we call r the setpoint value. Without additional knowledge, the value of r ought to be adjusted (or defined per mode) to fit a performance-maximizing criterion in all useful situations that the AO would face and across the complete range of the effective modal S/N. The automatic driving of G_i to satisfy Eq. (12) is of course particularly convenient if and only if a unique or a small number of setpoint values may be found.

The δk_{crit} correlation value and the value r we target to lock it onto are empirically representative of the spectral energy ratio between the low-frequency atmospheric rejection residual and the overshoot peak near f_{crit} . This is shown in Figs. 3 and 4.

Generally speaking, using higher values for r imposes a more cautious and robust steady-state control solution with a lower loop gain, with a strong, near unit correlation at δk_{crit} separation. Lower values lead to more aggressive loop behaviors that might reach nearly divergent transfer functions, but with a maximized rejection of the low-frequency components. As $\alpha_i G_i$ approaches g_{crit} , $m_i^*[\delta k_{\text{crit}}]$ goes toward -1 , and the output of the system is a slowly dampened sinusoid of period $2\delta k_{\text{crit}}$ frames.

Furthermore, with the condition of Eq. (12) satisfied, CLOSE enforces a transfer function constraint that is independent of the sensitivity reduction α_i of the WFS. In other words, it automatically compensates for the optical gain effect using modal compensation coefficients.

Fig. 5. For effective S/N $\sigma_i = 1, 10$, and 100 , closed-loop residual variances (top panels) and resulting δk_{crit} autocorrelation of the measurements (bottom panels) as a function of the loop gain $\alpha_i G_i$ from 0 up to g_{crit} . The latency is taken as $\tau = 2T$. Black lines show the total residual variance $\int_f |\hat{\phi}_{\text{Res}}^i(f)|^2 df$ normalized to 1 at its minimum, blue represents the variance induced by turbulence $\int_f |\hat{h}_{\text{Atm} \rightarrow \text{Res}}^i \cdot \hat{\phi}_{\text{Atm}}(f)|^2 df$, and green shows the variance induced by noise $\int_f |\hat{h}_{\text{n} \rightarrow \text{Res}}^i \cdot \frac{1}{\sigma_i}|^2 df$.

3.3. Correlation-locking and minimum variance: Semianalytical solutions

We have introduced the concept of a correlation-locking condition to adjust the controller gains and achieve a given control transfer function independently of the sensitivity reduction. We propose to compare the solutions obeying Eq. (12) to a minimum variance (MV) modal integrator, verifying Eq. (2).

For the example developed in Figs. 3 and 4, the gain $\alpha_i G_i$ minimizing the variance of the wavefront residual is $g_{\text{MV}} \approx 0.231$. At this gain value, we obtain $m_i^*[\delta k_{\text{crit}}] = -0.025$. This numerical observation indicates that $r = 0$ should be pursued as a candidate setpoint.

The computations are easily generalized for values other than the sensitivity-corrected S/N $\sigma_i = 10$ used for the examples. We show in Fig. 5 the variations of $m_i^*[\delta k_{\text{crit}}]$ and of the residual variance (Eq. (2)) with $\alpha_i G_i$, for S/N values $\sigma_i = 1, 10$, and 100 . We also plot the two components of the variance that are due to the turbulence residual and the noise propagation, that is, the two components that the CLOSE servo-loop seeks to balance optimally. The observed variations of $m_i^*[\delta k_{\text{crit}}]$ confirm its monotonicity with the gain as well as the relationship inferred in Sect. 3.2: a negative observable $m_i^*[\delta k_{\text{crit}}]$ (or setpoint r) relates to a high gain loop, and a positive $m_i^*[\delta k_{\text{crit}}]$ to a low gain loop. For all three S/Ns, Fig. 5 show the approximate match between the gain yielding minimum variance and the intercept $m_i^*[\delta k_{\text{crit}}] = 0$. This further establishes that $r = 0$ should be pursued as a special value that can be used in a variety of situations.

Generalizing these analyses further, we show in Fig. 6 the values of $\alpha_i G_i$ resulting in setpoints r of $-0.5, -0.1, 0, 0.1$, and 0.5 , together with the MV solution for effective S/N values $10^{-1} - 10^4$. The gain g_{MV} minimizing the variance for a given S/N is found following Gendron & Léna (1994) by numerically minimizing the joint Eqs. (4) and (5). CLOSE solutions are found by numerically solving $m_i^*[\delta k_{\text{crit}}] = r$ for $\alpha_i G_i$ through Eqs. (6) and (7). Additional data for $\tau = 0$ and $\tau = T$ are presented in Appendix B.

First, we note that the correlation-locked solution for an empirical $r = 0$ is a remarkably close match to the MV solution g_{MV} . The discrepancy at worst reaches 20% and 15% of g_{MV} in the two knees of the curve, near $\sigma_i = 3$ and $\sigma_i = 100$, respectively. We also confirm the statement made earlier:

Fig. 6. Minimum variance gain g_{MV} and solutions found using CLOSE with five different setpoint values, depending on the sensitivity-corrected S/N σ_i . These computations are performed using the modal input spectrum of Eq. (1) with $f_i = 1$ Hz.

Positive setpoints lead to lower values of $\alpha_i G_i$ for a given σ_i , and negative ones lead to a higher gain value. Interestingly, negative setpoints impose a gain floor even when the S/N is extremely poor. To achieve $r < 0$ over essentially white noise, the controller forcefully increases the gain so as to introduce a detectable oscillation of the amplified noise. Overall, $r = 0$ seems to provide an approximate near-MV solution over most of the σ_i range, with some room for improvement using a combination of $r > 0$ values for the higher end of the S/N range. The exact discrepancy between the g_{MV} and the CLOSE solution at null setpoint of course depends on the actual spectrum of turbulence, in particular, on the cutoff frequency f_i .

These semianalytical simulations altogether show that reaching the CLOSE steady-state solution using $r = 0$ can provide a near minimum variance solution throughout the entire range of σ_i . This makes it a relevant control technique for all the controlled DM modes for any choice of guide-star magnitude, for any amount of input turbulence, and for any sensitivity reduction α_i because all these parameters merely factor in the computation of σ_i . With a unique setpoint value, semianalytical computations offer a simple criterion that would enable fully automatic (almost) minimum variance integral control even for sensors with poorly modeled sensitivity variations.

We validate these claims using end-to-end numerical simulations in Sect. 5. We propose in the next section a practical implementation of the steady-state equations presented above.

4. Practical implementation

4.1. Real-time

Modal WFS telemetry $m_i[k]$ is obtained at each time-step k as the product of the WFS output by the modal reconstructor **Rec**. Two online autocovariance estimators are built from the $m_i[k]$ using discrete integrators,

$$\begin{aligned} N_i^0[k] &= p m_i[k]^2 + (1 - p) N_i^0[k - 1] \\ N_i^{\text{crit}}[k] &= p m_i[k] m_i[k - \delta k_{\text{crit}}] + (1 - p) N_i^{\text{crit}}[k - 1], \end{aligned} \quad (13)$$

where k is the time-step index and $p \in [0, 1]$ is a smoothing parameter. N_i^0 and N_i^{crit} are thus low-pass time-filtered estimates such that N_i^0 tracks the series variance $\text{Var}_k(m_i[k])$, and their ratio N_i^{crit}/N_i^0 tracks the critical autocorrelation $m_i^*[\delta k_{\text{crit}}]$. After an

empirical optimization of the parameter p , we opted for fast integrators with $p = 0.3$ for all simulations presented in Sect. 5. One objective achieved through this smoothing parameter is that the gap of δk_{crit} frames between the joint estimations of the variance and of $m_i^*[\delta k_{\text{crit}}]$ is bridged.

After the AC estimation, the modal gains G_i are updated using multiplicative increments as follows:

$$G_i[k] = G_i[k - 1] \times \left[1 + q^\pm \left(\frac{N_i^{\text{crit}}[k]}{N_i^0[k]} - r \right) \right]. \quad (14)$$

The r parameter is the loop setpoint as defined in Sect. 3.

The $q^\pm$ learning factor may encompass two different values, with either q^+ and q^- used depending on the sign of

$$\frac{N_i^{\text{crit}}[k]}{N_i^0[k]} - r,$$

q^+ being used for G_i increases in Eq. (14), and q^- for decreases. This asymmetry is kept as an option to make the algorithm more reactive to overshooting transients (using $q^- > q^+$) as compared to tracking gain increases due to a transfer function that is deemed too slow (using q^+).

This asymmetric tracking is not used for results shown in this paper, but was used in preliminary numerical simulations (and throughout Deo et al. 2019b). It allows maintaining stability at higher $q^\pm$ values. However, using $q^+ \neq q^-$ induces a bias in the mean value achieved for $m_i^*[\delta k_{\text{crit}}]$ that then differs from r . We empirically observed that for $q^- = 5q^+$, for example, $r \in [-0.2, -0.1]$ should be used to retrieve the performance seen with $r = 0$ in the $q^+ = q^-$ case.

The $q^\pm$ learning factors determine the time constants associated with the convergence and tracking ability of the CLOSE servo-loop. We infer that for a real AO system, $q^\pm$ values in the range of 10^{-3} – 10^{-4} should be used (assuming 500 Hz frequency), hence providing typical time constants in the 2–20 s range. The ideal choice of $q^\pm$ will probably remain dependent on the system and will certainly require some adjustments accounting for robustness and responsiveness to variations of turbulence conditions, vibrations, or other transient events. While the theoretical derivations were most accurate using $r = 0$, we do not exclude that for each system, some sort of empirical tweaking of r may be necessary to either privilege consistent stability or aggressive rejection.

4.2. Computational strain

Implementing CLOSE in a real-time fashion is of course expected to increase the AO loop computational requirements. While the AC estimations and gain updates themselves (Eqs. (13) and (14)) are negligible compared to the matrix-vector multiplication (MVM), having the $m_i[k]$ available in real-time requires to perform the reconstruction in two successive MVM steps. The first MVM converts WFS measurements into modal values, with a computational burden nearly identical to the usual overall MVM from measurements to DM commands. The second step computes DM increments from modal values, with a nearly square matrix with a size of the number of actuators. While this two-step operation is not universal, it is worth noting that it is used routinely on some instruments (e.g., Guyon et al. 2018). While this two-step technique may become a computational bottleneck in particular for high-contrast systems on ELTs, it is currently baselined for the control of the deformable M4 on ESO's ELT (Bonnet et al. 2019).

For a typical PWFS AO system with some spatial oversampling (the pixel projected size is smaller than the DM pitch), the number of pixels that is read out is typically 5–6 times the number of actuators because there are four pupil-like images, times the square of the oversampling factor. The number of outputs of the PWFS is therefore 2.5–6 times the number of modes, depending on whether slopes-map preprocessing is used or not. The algorithmic interest and computing cost of skipping the preprocessing is discussed in Deo et al. (2018b), for example. While the first MVM execution time depends on the WFS output dimensionality, the second MVM only depends on the number of controlled modes, with a smaller but not negligible computational footprint. As an example, we measured the real-time computer (RTC) computation time for the AO simulation setup we used in Sect. 5 (see Table 2): $881.6 \pm 4.2 \mu\text{s}$ using a single MVM, against $1035.6 \pm 3.6 \mu\text{s}$ using two cascaded MVMs. These timings were achieved using a single Nvidia Tesla P100 graphical processor.

4.3. Offline implementation

If the RTC software cannot be altered on an existing system or if the additional strain is not acceptable within the RTC specifications, CLOSE can be implemented in a block-wise manner. All estimators, gain updates, and command matrix updates are performed in offline time, certainly in another process, and preferably on another machine, over batches of contiguously recorded measurements. This buffered strategy enables deploying CLOSE on nearly any existing AO system that reliably provides its WFS telemetry without excessive delays.

A time-continuous buffer of K WFS measurements is forwarded to the CLOSE process, which turns them into modal measurements $m_i, \dots, m_i[K-1]$ using the modal reconstructor **Rec**. For each mode, the AC estimators of Eq. (13) are replaced by the direct computation of the normalized δk_{crit} -shifted AC term over the telemetry buffer,

$$N_i^{\text{block}} = \frac{1}{K - \delta k_{\text{crit}}} \frac{\sum_{k=\delta k_{\text{crit}}}^{K-1} m_i[k] m_i[k - \delta k_{\text{crit}}]}{\frac{1}{K} \sum_{k=0}^{K-1} m_i[k]^2}. \quad (15)$$

The gain-updating equation can then be performed,

$$G_i[\text{new}] = G_i[\text{previous}] \times \left[1 + q^\pm (N_i^{\text{block}} - r) \right], \quad (16)$$

using $q^\pm$ factors adjusted for the longer integration time and the increased SNR on AC estimation. Typically, $q^\pm$ ought to be larger by a factor $\sqrt{K}$ for a dynamical effect comparable to the real-time implementation. The new command matrix can then be computed accounting for the new G_i values, and when all side-tracked computations are completed, can be set into the RTC.

5. Numerical simulation results

This section describes some end-to-end numerical simulations that demonstrate the performance achieved with CLOSE when applied to the MICADO SCAO design (Clnet et al. 2019; Vidal et al. 2019). The main parameters of the system and the simulations are summarized in Table 2. All simulations were performed using the COMPASS platform (Ferreira et al. 2018b).

In Sect. 5.1 we analyze the convergence of modal gains when the AO loop is bootstrapped. In Sect. 5.2 we verify the steady-state performance achieved using $r = 0$ for various S/N levels.

Table 2. AO numerical simulation parameters.

Numerical simulation configuration	
Telescope	$D = 39$ m diameter ELT pupil model (no spider arms)
Turbulence	von K��rm��n, ground layer only r_0 at 500 nm: [8.9–21.5] cm $L_0 = 25$ m Speed: 10 m s^{-1} [10–40]
Guiding	On-axis natural guide star Zero point: $2.6 \times 10^{10} \text{ ph s}^{-1} \text{ m}^{-2}$ Magnitude M_R : [0–18]
DMs	Tip-tilt mirror Hexagonal M4 model pattern Pitch of 54 cm, coupling of 0.24 4302 controlled actuators
PWFS	
Subapertures	92×92 – pixel size 42 cm. 24 080 pixels used for control ^(a)
Wavelength	Monochromatic, 658 nm
Throughput	0.28 (including quantum efficiency)
Modulation	Circular, $4 \frac{\lambda}{D}$ radius
Readout noise	$0.3 e^-$
RTC controller	
Loop rate	$f_s = 500 \text{ Hz}$
Latency	$\tau = 2T = 4 \text{ ms}$ ^(b)
Method	Two step modal linear integrator Pixels $\rightarrow [m_i] \times [G_i] \rightarrow$ actuators
Basis	DM Karhunen–Lo��ve basis ^(c) Real-time implementation
CLOSE	$p = 0.3$; $q^\pm = 10^{-3}$; $\delta k_{\text{crit}} = 5$ $r = 0$. [–0.5, 0.5].

Notes. Values within brackets indicate probing ranges for various simulations reported across Sect. 5. Values preceding brackets indicate the default value for the parameter. ^(a)Bypassing the slope maps computation according to Guyon (2005), Clergeon (2014), and Deo et al. (2018b). ^(b)Extended to τ in $[0, T, 2T]$ in Appendix C. ^(c)Ferreira et al. (2018a).

This is expanded in Sect. 5.3 by exploring the effect of different setpoint values. Finally, we show in Sect. 5.4 the dynamic behavior of CLOSE when it is exposed to sudden changes in conditions. Throughout the following sections, seeing conditions are referred to using the Fried parameter r_0 , always given at 500 nm; and guide-star brightnesses are identified by the R -band magnitude (M_R), related to the photon flux per the zeropoint and system throughput given in Table 2. AO performance is most often measured in terms of H -band long-exposure (LE) Strehl ratios (SR), computed from simulated monochromatic point-spread functions at 1650 nm. SR comparisons given in percent are always in percentage point units, never relative variations.

5.1. Gain convergence upon closing the loop

We first investigate the dynamics of the modal gains upon closing the AO loop with CLOSE enabled. These simulations were all initialized identically, regardless of r_0 or M_R . We opted for the starting value $G_i[k=0] = 0.5$ for all modes. With the system latency $\tau = 2T$ simulated, the critical gain value was $g_{\text{crit}} \approx 0.61$. Because the sensitivity reduction α_i is always smaller than 1, this ensured that the loop was closed with a comfortable

Fig. 7. Convergence of CLOSE gains on the 2.0 s following closing of the AO loop for guide stars of $M_R = 0$ and 16 and atmospheric r_0 of 14.5 and 9.0 cm. All G_i are initialized to 0.5 at the start (blue line). G_i values are shown as averaged over the time windows given in the legend. Curves are smoothed along the i index for clarity. Final SRs are given in H band and are computed from the cumulative exposure over the last 200 ms ($k \in [900, 1000]$).

stability margin. From these initial 0.5 values, the G_i were driven by CLOSE to their steady-state values, accounting both for non-linearity compensation and temporal variance minimization.

We show in Fig. 7 the temporal averages of the 4301 modal gains G_i for several time windows within the first two seconds after the AO loop is closed. These simulations were performed for four different cases, with r_0 of 14.5 and 9.0 cm, and guide stars with brightnesses of $M_R = 0$ and $M_R = 16$.

For the bright cases, steady state is reached by frame $k \approx 500$, that is, within one second. The process is slightly slower for the $M_R = 16$ cases, with a continued convergence of the G_i between frames $k = 500$ and $k = 1000$. Simulations at $M_R = 0$ are essentially performed with an infinite S/N. The dynamical gains $\alpha_i G_i$ therefore evolve nearby the maximum stability value g_{crit} , and the G_i coefficients reached in steady state essentially reflect the inverse α_i^{-1} of the PWFS sensitivity reduction. The G_i curves reached at the end of convergence in $M_R = 0$ cases are in good accordance with the abacuses presented in previous work (Deo et al. 2018a, 2019a) that were obtained by directly measuring α_i sensitivity reductions on static turbulence screens, with α_i typically depending on mode index i as decreasing up to mode 30, which contains spatial frequencies corresponding to the modulation radius, then again increasing roughly as a power law up to the highest-order modes.

Some more insight into the bootstrapping of CLOSE can be gained by inspecting temporal series of modal gains, as shown in Fig. 8. These series correspond to four modes and are the same data set as in Fig. 7 for $r_0 = 14.5$ cm, $M_R = 0$. While the AO loop bootstraps in only a few frames, the convergence of CLOSE takes longer and induces some modal gain oscillations after the initial ramp-up. The amplitude and time constants of these oscillations certainly depend on a number of parameters, and impor-

Fig. 8. Time series of CLOSE gains for four select modes in the 2.0 s following closing of the AO loop for the simulation $r_0 = 14.5$ cm, $M_R = 0$ shown in Fig. 7.

tantly on the learning factor $q^\pm$. In Fig. 8, various periods are observed, from ≈ 1 s down to smaller oscillations typically every ~ 30 frames (60 ms).

The randomness of the turbulence screens is certainly a factor in continuous variations of G_i , and the oscillations continue even as the AO integrator and CLOSE reach steady states, beyond 1–2 s. of runtime, but with little effect on the correction quality. Figures 11 and 12 show complementary data to those in Figs. 7 and 8, including longer time series and higher Strehl ratios.

Furthermore, the physical nature of the PWFS is certainly a contributor to the G_i fluctuations as compared to a purely linear sensor. When G_i reaches values that are too high, or if we had started from G_i values higher than the steady state that is ultimately reached, the optical gain phenomenon itself helps to maintain stable suboptimal control states while the controller performs the slow gain decrease. To summarize this effect, reduced atmospheric residuals (e.g., bootstrapping or improving conditions) induce an increase in WFS sensitivity. When the gains in G_i are overset, transfer functions become highly resonant or temporarily unstable. In turn, the added wavefront residual from transitory divergence of the loop or noise oscillations induces a reduction in the PWFS sensitivity, however. This regime is progressively stabilized as the adaptive filtering eventually adapts the command law to the ongoing conditions and reaches a steady-state regime. These transitory effects are further discussed in Sect. 5.4.

When we compare in Fig. 7 the behavior in the $M_R = 0$ and $M_R = 16$ cases, we observe the effect of the implicit optimization of the transfer function, with steady-state gain values dampened by typically 20–50% in the dim case relative to the bright one, depending on the mode number and r_0 . The results presented in Fig. 7 tend to validate that without any priors and regardless of the PWFS sensitivity reduction, CLOSE successfully drives the modal integrator gains to convergence in a period of 1–2 s.

5.2. Performance in stationary conditions for $r = 0$

In addition to the adaptive capability of CLOSE shown in Sect. 5.1, we investigated the steady-state AO performance. In order to perform this analysis, we generalized similar simulations as performed in Sect. 5.1 to a wider range of r_0 values (based on statistics provided by ESO within the frame of the development of ELT instruments) and guide-star magnitudes (11.5–17.5).

The measured performances are shown in Fig. 9 (top), with the long-exposure H -band SR plotted against the star magnitude and computed for five different seeing conditions. For all results

Fig. 9. Long-exposure SR in H band obtained for end-to-end simulations of the MICADO SCAO setup with CLOSE (*top*), and with a manually optimized scalar integrator (*bottom*) for guide-star magnitudes $M_R = 11.5$ – 17.5 and seeing conditions $r_0 = 8.9$ – 21.5 cm.

the SR was averaged over 2 s exposures, starting 3 s after the AO loop was closed, and with initial modal gains $G_i = 0.5$ as previously. These simulations were also performed with a manually optimized scalar integrator ($G_i = \text{constant}$, with one constant value per seeing and magnitude), and the performance is shown in Fig. 9 (*bottom*). This again shows the improvement that is achieved using fine-tuned modal control, which was demonstrated in previous research. In particular, using modal control enables us to (1) improve the performance in poor seeing, even with bright guide stars, because nonlinearity is then a dominant member of the error budget. It also enables us (2) to increase the performance at the faint end, with a gain of typically +0.5 mag for an identical objective.

While the simulation setup is not a perfect simulation of the MICADO system, Fig. 9 (*top*) is indicative of the performance that could be achieved using CLOSE on such a system. If we exclude the worst seeing conditions $r_0 = 8.9$ cm, a performance of 70% SR or better is always achieved for bright stars, it even reaches 89% in excellent seeings. 40% or better is achieved for M_R up to 16. SRs better than 40% are also achieved for bright stars for $r_0 = 8.9$ cm. This potentially enables us to salvage poor-seeing nights for a number of scientific cases with milder correction requirements. For the metric used here, CLOSE performs equally to other previously introduced modal compensation techniques for the PWFS (Korkiakoski et al. 2008b; Deo et al. 2019a; Chambouleyron et al. 2020), but with the added value of automation and adaptability, without the offline computations, seeing estimations, dithering signals, or modifications to the optical setup required by these methods.

5.3. Validation of the setpoint

We have shown throughout Sects. 3 and 4 that CLOSE is entirely configured with very few parameters, namely, δk_{crit} , which ought to be chosen depending on the system latency, and p , $q^\pm$, and

r being the choice of the operator. While the temporal filtering introduced with p and $q^\pm$ is easily apprehended and only affects the performance during transitory regimes, the determination of r is subject to more caution because it determines the final performance. Although semianalytical computations indicated $r = 0$ as an apparently universal choice, these computations were the conclusion of a number of modeling hypotheses and approximations, as described in Sect. 2.

Our objective here is to validate whether $r = 0$ remains an appropriate choice in most situations. We performed end-to-end simulations extending those presented in Sect. 5.2, now considering various r setpoints from -0.5 to 0.5 and wind speeds from 10 to 40 m s^{-1} . The performances measured are shown in Fig. 10. Some minor discrepancies can be found between Figs. 9 and 10, which are explained by the different durations of simulated long exposures and the mismatched random turbulence screens. The SRs for Fig. 10 were computed on 600 ms exposures, following a 600 ms bootstrap for CLOSE and the AO. As compared to Fig. 9 ($2 + 2$ s), this was a necessary speedup given the large number of numerical simulations that were performed. We also note that LE SRs are generally determined to no better than 2–3% of standard deviation over the distribution of turbulence screens.

In addition to the expected variations in SR with seeing, guide-star magnitude, and wind speed, we observe in Fig. 10 that the setpoint yielding the maximum SR, r_{max} thereafter, is almost always -0.1 or 0 , except for two cases ($M_R = 12$, $r_0 = 12.8$ and 8.9 cm and 30 m s^{-1}) where $r_{\text{max}} = -0.2$. When for a given r_0 , M_R , and wind speed we obtain $r_{\text{max}} \neq 0$, we observe that the difference with the corresponding performance at $r = 0$ is generally only 1–2%. The slight bias toward $r_{\text{max}} < 0$ may partly be explained by the reduced simulated time for CLOSE convergence. Cases with lower r would therefore have increased the modal gains more effectively in the allocated 300 frames because of the proportional effect introduced in Eq. (14).

The variations in SR with r are clearly determined from Fig. 10 as an asymmetric bell curve with a longer decrease on the $r > 0$ side. With r larger than the optimal value, the AO uses slower modal integrators with less turbulence rejection, which are more forgiving with regard to performance than r being too small. The latter cases introduce a buildup of noise that ultimately leads to diverging oscillations as r becomes too negative, hence a faster decrease in SR for simulations with r smaller than the optimum.

A few outlying cases show a significant performance gap between the maximum SR and the SR achieved for $r = 0$, up to a greatest difference of 6.3% ($M_R = 15$, $r_0 = 12.8$ cm, 20 m s^{-1}). These cases, where $r = 0$ induces a noticeably suboptimal performance, are all found for wind speeds of 20 m s^{-1} or more. With the high latency of $\tau = 2$ frames simulated here, the higher wind speed induces a narrow acceptable range of gain $\alpha_i G_i$ for each mode to achieve near-optimal rejection, and thus a narrow range of r values leading to this optimization with CLOSE. As discussed in Sect. 3.3, while the CLOSE solution achieved for the modal gain is an empirical close match to the optimum rejection solution, we may be reaching the limits of this approximation for cases with high latency and high wind speeds.

To investigate the usability of CLOSE with $r \approx 0$ even further, we also performed similar simulations for different values of the system latency using $\tau = 0$, 1 or 2 frames and using AO setups other than the MICADO SCAO, namely two SCAOs on an 8 m telescope, using a PWFS and a Shack–Hartmann (SH) WFS, respectively. The configuration of these simulations and the obtained results are shown in Appendix C. Beyond the

Fig. 10. Long-exposure SR in H band (color and text) for stationary simulations exploring different conditions regarding the seeing (*outer rows*), the guide-star magnitude (*outer columns*), the wind speed (*inner rows*), and the CLOSE setpoint (*inner columns*) for the MICADO SCAO simulation. Color scales are local to each of the subplots and do not match each other.

results shown in Fig. 10, the extensive simulations exposed in Appendix C confirm that the setpoint $r = 0$ may be used as a baseline for a variety of AO systems in a large number of observation situations. The performance and steady-state behavior may eventually be fine-tuned by adjusting the r setpoint upon empirical criteria if deemed necessary.

5.4. Transients in observation conditions

Beyond the performances of correlation-locking observed in stationary conditions exposed in Sects. 5.2 and 5.3, we propose here to analyze the dynamical capabilities of CLOSE in situations where the PWFS sensitivity or the illumination vary rapidly. For this purpose, we designed two simulations with evolving conditions. The first simulation emulates a seeing burst in which atmospheric conditions degrade dramatically during a short period, and the second simulation emulates a passing cirrus, with an equivalent drop of 3 mag (-94% of flux) of the guide star.

The results for these simulations are shown in Figs. 11 and 12. We proceeded identically for both transients, simulating 8 s (4000 frames) of total runtime decomposed as 2 s of bootstrapping the AO loop from flat gains $G_i(t = 0) = 0.5$, followed by 4 s in degraded conditions, and finally 2 s after reverting to the initial conditions. Both Figs. 11 and 12 show the time series of the H -band SR and of the modal gains G_i for four select modes. The lower panels in Figs. 11 and 12 show a snapshot of the modal gain vector $\mathbf{G}$ over all 4301 DM modes at times $t = 2, 6$, and 8 s, that is, immediately before the changes in seeing or brightness, and at the end of the simulation.

Fig. 11. *Top:* time series of the H -band SR and modal gains G_i for a few select modes during the simulation of a burst of seeing; $r_0 = 21.5$ cm for $t < 2$ s and $t > 6$ s and $r_0 = 8.9$ cm for $2 < t < 6$ s. Curves are smoothed using a 10 ms window. *Bottom:* snapshot of the modal gain vector $\mathbf{G}$ at time $t = 2, 6$, and 8 s.

Seeing burst. For the seeing burst (Fig. 11), the turbulence screen was scaled up between $t = 2$ and 6 s to simulate a r_0 change from 21.5 cm to 8.9 cm while keeping a guide star of brightness $M_R = 14$. The sudden changes in r_0 at $t = 2$ and 6 s induce discontinuities in the wavefront because we performed an instantaneous scaling across the entire aperture. This induced very short drops to 0% SR (a few frames). At $t = 2$ s, the modal gains optimized for $r_0 = 21.5$ cm do not permit reaching optimal

Fig. 12. Similar to Fig. 11. Time series and modal gain snapshots during a cloud-like transient, simulated as a guide-star magnitude change: $M_R = 14$ for $t < 2$ s and $t > 6$ s and $M_R = 17$ for $2 < t < 6$ s.

performance immediately now that $r_0 = 8.9$ cm. For $2 < t < 3$ s, CLOSE drives the G_i up by a factor 2–4 to accommodate the OG-induced sensitivity reduction, which is significantly stronger at $r_0 = 8.9$ cm. This second long convergence to the new optimal state allows an SR improvement up to the 30–35% consistently measured for $3 < t < 6$ s, which is the expected performance in these stationary conditions ($r_0 = 8.9$ cm, $M_R = 14$; as seen in Fig. 9).

Similarly, the second transition back to $r_0 = 21.5$ cm at $t = 6$ s shows an immediate SR performance about 45% below the expectation, indicating that the system is out of tune because the G_i are now beyond reasonable values for $r_0 = 21.5$ cm. The modal gains are automatically decreased, however, and the SR converges to 85% within 0.5 s. This is consistent with the magnitude of the time constants involved because we used a learning factor $q^\pm = 10^{-3}$ at a frame rate of 500 Hz.

Photometric variations. The “cirrus simulation” (Fig. 12) was run using the median seeing $r_0 = 14.5$ cm, with a change of guide-star brightness from the initial $M_R = 14$ down to $M_R = 17$ during the perturbation interval $2 < t < 6$ s. While such a flux attenuation qualitatively appears as a major change, the effect on the gain is rather mild when compared to the previous case. The reaction of CLOSE that can be observed in Fig. 12 is that an appropriate overall reduction in modal gains is progressively introduced during the transient and then reverted after $t = 6$ s. This behavior is typically expected from modal gain optimization with S/N. The SR temporal evolution shows that the progressive optimization of the transfer functions for $M_R = 17$ allows for an improvement from 10 to 20% of SR within the 4 seconds spent at this brightness level.

It is worth noting that while the modal gain curve reached at $t = 6$ s only shows a modest reduction compared to $t = 2$ s, this reduction is in fact the product of two counterbalancing effects: First, a reduction of $\alpha_i G_i$ to improve the noise rejection, but coupled to an increase in the sensitivity compensation α_i^{-1} because the mean residual wavefront is notably increased, which induces a stronger OG sensitivity reduction in the PWFS response. These two effects are entirely implicit: As the flux drops by 3 mag, the corrective capability of the AO immediately decreases, the residual wavefront increases, and thus the sensitivity α_i decreases as well. Factored in the transfer function, the gain $\alpha_i G_i$ is decreased

immediately and compensates for the sudden shift in S/N. Afterward, CLOSE accounts for the fine-tuning of the G_i to more precisely set the rejection at the new sensitivity level. The opposite occurs at $t = 6$ s: The sharpening of the correction improves the sensitivity, and the dynamical gain $\alpha_i G_i$ immediately increases and provides more turbulence rejection, now with less noise to amplify. This effects explains why compared to the seeing burst, (1) the fluctuations in G_i are significantly milder and (2) the return-to-steady-state time is much shorter.

Finally, as with the seeing burst, we note that the performance in the final state $t > 6$ s is identical to the beginning $t < 2$ s. This indicates that no diverging signals were accumulated during the transitions and that the AO is capable of reverting cleanly after such transients.

Synthesis. From the two experiments performed with sudden changes in observing conditions, we may conclude that CLOSE efficiently and automatically accommodates sudden significant changes in S/N, whether from the seeing or the guide-star brightness. This robustness to transients is observed for improving or deteriorating conditions. This efficient convergence toward nominal performance, when the seeing improves in particular, underlines the robustness of the proposed AO control scheme. When the seeing suddenly improves, the AO is configured with gains that are significantly too high for the newly set turbulent conditions, but the control scheme allows reoptimizing the gains without the loop diverging.

6. Discussion

6.1. In general: An interaction matrix transform

In the most general sense, as developed in Sects. 2 and 3, CLOSE is an automatic minimum variance controller for servo-loops whose input is built of a low-frequency rumble and a white noise floor (Eq. (1), Fig. 3). When these conditions are met, as is the case with astronomical adaptive optics, and when independent modal loops can be decoupled, CLOSE is an excellent tool for optimizing the system regardless of miscalibrated sensitivity factors.

CLOSE is therefore adequate for any system that admits a description as in Fig. 1, where an eigenbasis can be found that properly describes the nonlinearity of the system, that is, diagonalizes (well enough) the modal confusion matrix $\mathbf{A}$. This is the case for the PWFS, as has been proven experimentally (Deo et al. 2018a) and analytically (Fauvarque et al. 2019; Chambouleyron et al. 2020). This was extended to the entire class of Fourier-based wavefront sensors.

Furthermore, it has been demonstrated that a number of configuration changes of PWFSs are very well described by variations in the modal matrix $\mathbf{A}$, while the interaction matrix $\mathbf{dWFS}$ can be kept to the unresolved, bright, on-axis, single-star matrix acquired in the laboratory (or simulated). These changes include online changes to the modulation radius, or tracking double stars of various separations and orientations relative to the pyramidal prism edges, as well as other extended objects (Titan, uncompensated atmospheric dispersion beacon; Vidal et al. 2019). For binary stars, it has in particular been shown that using CLOSE restores the performance of the servo-loop to near single-star levels, without the additional burden of biasing the loop setpoint to drive one binary component to the PWFS apex, which is the usual and nonetheless efficient alternative.

All of these transformations, similarly to optical gain, can be approximated by Fourier filters that alter the description of

the sensor and thus are adequately tracked by CLOSE. Ideally, a Fourier basis would have to be used to properly diagonalize $\mathbf{A}$. Our experience and numerical analyses demonstrate that a frequency-ordered KL basis is a sufficient approximation. We previously studied (Deo et al. 2018a) for the case of optical gain alone that the scaling applied by $\mathbf{A}$ on Fourier modes of identical norm is isotropic within a few percent. This proves that KL modes are adequate for such an operation.

6.2. Further required steps

Finally, we would like to point out a few elements that are currently beyond the scope of what CLOSE is capable of achieving. One of the most prominent issues with high-order PWFSs is the proper subtraction of non-common path aberrations (NCPAs) as the servo-loop setpoint, which issue is tightly coupled with the estimation of modal OG compensation coefficients. In the state presented in this paper, CLOSE does not enable an independent estimation of the transfer function gain $\alpha_i G_i$ and the OG compensation coefficient α_i^{-1} , as we observed in the cloud transient analyzed in Sect. 5.4. Obtaining α_i^{-1} is necessary for a proper OG-adjusted compensation of the setpoint. Other techniques, which include injecting probe signals into the system DM, allow an explicit retrieval of α_i^{-1} (Esposito et al. 2020). While CLOSE does not include this feature, it is very compatible with such techniques but is much more reactive to short-term transients. A joint implementation of several modal gain estimation techniques would be extremely beneficial for an all-in-one solution for PWFSs control, including minimum variance laws, proper NCPA compensation, shift and misalignment tracking (Héritier et al. 2018), etc.

One additional comment should be made about the potential limitations of the modeling described in Sect. 2. The proper variance minimization obtained with CLOSE relies on the clear separation of two frequency components from the turbulence and the noise in the resulting measurement PSD: $|\hat{m}_i|^2(f)$. The requirements of this paper exceed this by requiring that the modal noise is white and that the turbulence spectrum obeys Eq. (1). We proved with the Shack–Hartmann results shown in Appendix C, however, that the decay exponent $-17/3$ of Eq. (1) may allow some tolerance because it is then actually $-11/3$ due to additional subaperture aliasing. Together with using different von Kármán outer scales L_0 , this is very likely to affect the quantitative conclusions of the semianalytical derivations described in Sect. 3.3, but we are confident that it will not compromise the monotonicity of the relationship between the δk_{crit} correlation and the loop gain $\alpha_i G_i$. Thus, a simple reassessment of the loop setpoint r to try and achieve quasi-optimality is likely to be sufficient for situations like this.

Telescope vibrations are a less trivial modification of the input turbulent spectrum. Naturally, CLOSE cannot achieve any better vibration rejection than other nonpredictive integrators. As for other spectrum modifications, a reassessment of r could help in achieving near optimality, as good as this may be in the presence of narrow vibration peaks. Low-frequency (e.g., <10 Hz) vibrations would implicitly be registered as a degradation in seeing, and would require biasing the setpoint toward $r > 0$ to prevent the corresponding gain increase. Conversely, vibrations between $f_{\text{crit}}/2$ and $3f_{\text{crit}}/2$ would be registered as noise energy and require biasing the setpoint toward $r < 0$.

We foresee that vibrations at or near f_{crit} will cause specific issues. All of our simulations have shown that CLOSE, seeking its near-MV solutions, tends to push the loop gain very close to its critical value when the S/N is good enough. This could lead

to transients with very high vibration amplification, for instance, as the vibration restarts after a pause during which $\alpha_i G_i$ was boosted near g_{crit} . These situations will require the implementation of specific fail-safes to avoid the sudden divergence of the AO loop. More generally, CLOSE in the presented version does not currently cope well with spurious divergences of the AO loop: The loss of WFS feedback associated with DM divergence (or clipping at maximum value) will induce CLOSE (if $r \leq 0$) to continue increasing gains toward infinity, even though the loop has already diverged. It is now necessary to experiment with various modifications of the input spectra (correlated noise, vibrations, non von Kármán spectra) and the integral command law (leaky integrators, linear predictive methods) to investigate the extent to which and the versatility with which we may obtain satisfactory results with the version of CLOSE presented in this paper, and which further improvements are necessary.

7. Conclusion

We have presented a self-regulating adaptive filtering method that enables tracking the modal gains of the integrator controller in an AO system in hard or soft real-time. This technique allows keeping the system always close to peak performance against variable conditions (seeing, wind speed, and S/N) and the variable sensitivity of the WFS. Because it is entirely automatic and independent of WFS sensitivity fluctuations, the proposed scheme is particularly well suited for high-order PWFS designs on ELTs, which are prone to strong optical gain effects. This OG effect is well decoupled and stable against seeing statistics using a KL control basis of the DM as we use, however, thus enabling a separable modal analysis throughout.

Our method is called CLOSE, for correlation-locking optimization scheme. It leverages the monotonic relationship between the hidden actual transfer function gain and the loop resonance peak, whose amplitude is indirectly assessed through the temporal autocorrelation of the modal phase residuals computed at a single given time-shift δk_{crit} . The appropriate counter-reaction is then applied on the modal gain in order to lock the previously mentioned δk_{crit} correlation value to a chosen setpoint. Using semianalytical computations, we have demonstrated that for given conditions, there is an unequivocal relation between the transfer function gain and the value of the δk_{crit} correlation, and that correlation locking easily permits almost achieving a modal minimum variance criterion regardless of the modal S/N.

We have used extensive end-to-end AO simulations to show the versatility of CLOSE. Using the unique setpoint $r = 0$, we showed automatic convergence of the modal gains in a wide variety of simulated conditions using different wind speeds, guide-star magnitudes, and seeing parameters. We also demonstrated the automatic adaptability of the scheme to abrupt changes in observation conditions, hopping between different sets of optimal control gains in a matter of seconds. These simulations use the design parameters of the PWFS of the MICADO SCAO system, and compared to recent design studies for this system (Vidal et al. 2019), achieve nominal performance accounting both for optical gain modal compensation and modal transfer function optimization. This performance is achieved while entirely alleviating from extensive situation-dependent optimizations, database look-ups, or manual tune-ups, thus offering CLOSE as an all-in-one baseline strategy for AO modal control.

While CLOSE elegantly paves the way for solving variance minimization issues with nonlinear sensors, we cautioned

in Sect. 6.2 that the scheme in the version described in this paper does not provide an explicit estimation of the sensitivity reduction coefficient nor a fully fledged solution to the NCPA subtraction issue for sensors of varying sensitivity.

Work is currently underway to refine and expand the work described in this paper to make it usable for on-sky operations. Importantly, we will design and deploy the necessary fail-safes in case CLOSE suffers from a loss of reliable feedback from the modal outputs, or if the AO falls into nonlinear measurement pits. In these cases, the modal gain update equations tend to increase the gains exponentially. These situations must be automatically identified and unambiguously distinguished from transient sensitivity reductions. If its capabilities, convenience, and versatility are confirmed, we envision CLOSE as a core control technique for future AO systems, embedded within a larger framework of leveraging real-time telemetry for AO control.

Acknowledgements. This research is performed in the frame of the development of MICADO, first light instrument of the ELT (ESO), with the support of ESO, INSU/CNRS and Observatoire de Paris. V. Deo is supported by NASA grant 80NSSC19K0336. Additional support was provided through the WOLF project, ANR-18-CE31-0018 of the French National Research Agency (ANR), and the OPTICON Program of the European Commission (H2020, Grant number 730890). The authors thank the COMPASS development and support team for their continued support.

References

- Bonnet, H., Spyromilio, J., Véronaud, C., et al. 2019, in *ELT Wavefront Control Strategy*, 6th AO4ELT conference
- Chambouleyron, V., Fauvarque, O., Janin-Potiron, P., et al. 2020, *A&A*, **644**, A6
- Clénet, Y., Buey, T., Gendron, É., et al. 2019, in *MICADO-MAORY SCAO Preliminary Design, Development Plan and Calibration Strategies*, 6th AO4ELT Conference
- Clergeon, C. 2014, PhD Thesis, Observatoire de Paris - Subaru NAOJ
- Conan, J.-M., Rousset, G., & Madec, P.-Y. 1995, *J. Opt. Soc. Am. A*, **12**, 1559
- Costa, J. B. 2005, *Appl. Opt.*, **44**, 60
- Dai, G.-M. 1996, *J. Opt. Soc. Am. A*, **13**, 1218
- Davies, R., Alves, J., Clénet, Y., et al. 2018, *Proc. SPIE*, **10702**, 107021S
- Deo, V., Gendron, E., Rousset, G., et al. 2018a, *Proc. SPIE*, **10703**, 1070320
- Deo, V., Gendron, E., Rousset, G., et al. 2018b, *A&A*, **619**, A56
- Deo, V., Gendron, E., Rousset, G., et al. 2019a, *A&A*, **629**, A107
- Deo, V., Rozel, M., Bertrou-Cantou, A., et al. 2019b, in *CLOSE: a Self-regulating, Best-performance Tracker for Modal Integrator Based AO Loops*, 6th AO4ELT Conference
- Dessenne, C., Madec, P. Y., & Rousset, G. 1998, *Appl. Opt.*, **37**, 4623
- Esposito, S., Riccardi, A., Fini, L., et al. 2010, *Proc. SPIE*, **7736**, 773609
- Esposito, S., Pinna, E., Puglisi, A., et al. 2015, *Non common path aberration correction with non linear WFSs*, 4th AO4ELT Conference
- Esposito, S., Puglisi, A., Pinna, E., et al. 2020, *A&A*, **636**, A88
- Fanson, J., McCarthy, P., Bernstein, R., et al. 2018, *Proc. SPIE*, **10700**, 1070012
- Fauvarque, O., Janin-Potiron, P., Correia, C., et al. 2019, *J. Opt. Soc. Am. A*, **36**, 1241
- Ferreira, F., Gendron, É., Rousset, G., & Gratadour, D. 2018a, *A&A*, **616**, A102
- Ferreira, F., Gratadour, D., Sevin, A., & Doucet, N. 2018b, in *2018 International Conference on High Performance Computing Simulation (HPCS)*, 180
- Gendron, É. 1995, PhD Thesis, Univ. Paris Diderot, France
- Gendron, É., & Léna, P. 1994, *A&A*, **291**, 337
- Guyon, O. 2005, *ApJ*, **629**, 592
- Guyon, O., Martinache, F., Garrel, V., et al. 2010, *Proc. SPIE*, **7736**, 773624
- Guyon, O., Sevin, A., Gratadour, D., et al. 2018, *Proc. SPIE*, **10703**, 107031E
- Héritier, C. T., Esposito, S., Fusco, T., et al. 2018, *Proc. SPIE*, **10703**, 107034P
- Korkiakoski, V., Véronaud, C., & Le Louarn, M. 2008a, *Proc. SPIE*, **7015**, 701554
- Korkiakoski, V., Véronaud, C., & Le Louarn, M. 2008b, *Appl. Opt.*, **47**, 79
- Kulcsár, C., Raynaud, H.-F., Petit, C., Conan, J.-M., & de Leseqno, P. V. 2006, *Opt. Express*, **14**, 7464
- Liu, F., & Sanders, G. 2018, *Proc. SPIE*, **10700**, 1070013
- Madec, P.-Y. 1999, ed. F. Roddier, *Adaptive Optics in Astronomy* (Cambridge: Cambridge University Press), 131
- Montera, D. A., Brown, J. M., Buckman, M. D., et al. 2018, *Proc. SPIE*, **10703**, 107031H
- Ragazzoni, R. 1996, *J. Mod. Opt.*, **43**, 289
- Schatz, L. H., Males, J. R., Close, L. M., et al. 2018, *Proc. SPIE*, **10703**, 1070321
- Tamai, R., Koehler, B., Cerasuolo, M., et al. 2018, *Proc. SPIE*, **10700**, 1070014
- Vidal, F., Rozel, M., Deo, V., et al. 2019, in *Tests and Characterisations of the ALPAO 64x64 Deformable Mirror, the MICADO-MAORY SCAO AIT Facility*, 6th AO4ELT Conference

Appendix A: Expanding on AO transfer functions

We detail here the transfer functions of the AO loop as a complement to Sect. 2.3 and highlight the differences between a continuous-time description, as is necessary for a real physical system, and a discrete-time description that is adequate for AO simulators such as COMPASS. The material reviewed in Appendix A.1 is well known (e.g., Madec 1999; Kulcsár et al. 2006) but seldom presented, and we deem it necessary for a proper mathematical grounding of Sects. 2.3 and 3. Ultimately, this development lets us derive the values for f_{crit} and g_{crit} provided in Table 1.

A.1. Continuous- and discrete-time descriptions

Both the real continuous-time system (CTS) and the simulated discrete-time system (DTS) are clocked at the sampling period T , which is also the integration time of the WFS detector. In the DTS, the turbulence is not averaged over T , but sampled at a single instant. Likewise, the notion of the duration of action of the deformable mirror simply does not exist. Both systems, however, have in common the core of the control: they share the same numerical algorithms as are implemented in the real-time computer, which in both cases are numerical and in discrete time steps.

We introduce several transfer functions hereafter. The WFS performs an integration of the signal during T . In Fourier formalism, its transfer function can be written as

$$h_{\text{wfs}}(f) = \exp(-j\pi fT) \frac{\sin(\pi fT)}{\pi fT}, \quad (\text{A.1})$$

which is the Fourier transform of a rectangular window extending from $t = 0$ to $t = T$. We may see Eq. (A.1) as a smoothing by a zero-centered box: $\text{sinc}(\pi fT)$, combined with a half-frame delay, $\exp(-j\pi fT)$, thus ensuring h_{wfs} describes a causal system.

The WFS measurement is processed by the reconstructor matrix to be expressed in a modal space $m[k]$, we drop here the mode index i for clarity, and is integrated with a gain g into the modal command $v[k]$ using the classical discrete integrator,

$$v[k] = v[k-1] + g \cdot m[k], \quad (\text{A.2})$$

where k is the loop iteration number. The direct translation of Eq. (A.2) into the frequency space gives the expression of the transfer function of the discrete time integrator,

$$g \cdot h_{\text{digInt}}(f) = \frac{\hat{v}(f)}{\hat{m}(f)} = \frac{g}{1 - \exp(-2j\pi fT)}. \quad (\text{A.3})$$

It should be noted that $h_{\text{digInt}}(f)$ is nonphysical, first because it shows no computational time delay, and second because it is periodically undefined for every multiple of the sampling frequency $f_s = 1/T$: it is only appropriate for representing a discrete time system.

Now, in an real CTS, the RTC computation result $v[k]$ is sent to a digital-to-analog converter (DAC), which holds the (modal) command on the DM during the period T . The DAC, a zero-order hold, has the same transfer function as the WFS,

$$h_{\text{dac}}(f) = \frac{1 - \exp(-2j\pi fT)}{2j\pi fT}, \quad (\text{A.4})$$

also including a half-frame delay. With the association $h_{\text{wfs}} \cdot h_{\text{dac}}$, a full frame delay thus naturally appears.

Equation (A.4) joins the digital world back into the physical reality of the CTS by transforming the Dirac comb of numerical

samples into an analog signal. Noticeably, the association of the nonphysical numerical integrator with the DAC transfer function form the following product,

$$h_{\text{digInt}}(f) \cdot h_{\text{dac}}(f) = \frac{1}{2j\pi fT}, \quad (\text{A.5})$$

that is, the proper transfer function of a continuous-time integrator.

The final point of the discussion is that of the time delay τ . We defined the loop delay as the additional amount of time, on top of the causal process. For a CTS, this is the time spent between the end of the WFS integration and the beginning of the application of the command on the DM. In this case, the transfer function is

$$h_{\text{delay}}(f; \tau) = \exp(-2j\pi f\tau). \quad (\text{A.6})$$

When $\tau = 0$, the DM command $v[k]$ exactly applies during the acquisition of $m[k+1]$. This one-frame shift is built in the equations thanks to the two half-frames from the WFS and DAC. All the elements are introduced for the transfer function between the input turbulence ϕ_{Atm}^i and the modal measurements $m_i[k]$ in the case of a CTS,

$$h_{\text{corr,CTS}}(f; g) = \frac{1}{1 + g \cdot h_{\text{wfs}}(f) \cdot h_{\text{digInt}}(f) \cdot h_{\text{dac}}(f) \cdot h_{\text{delay}}(f; \tau)}. \quad (\text{A.7})$$

For the DTS, there is neither WFS nor DAC windowing, but we need to introduce an apparently artificial one-frame delay $\exp(-2j\pi fT)$ to account for the fact that a command computed at a given iteration can at best only serve for the next one,

$$h_{\text{corr,DTS}}(f; g) = \frac{1}{1 + g \cdot h_{\text{digInt}}(f) \cdot h_{\text{delay}}(f; \tau) \cdot \exp(-2j\pi fT)}. \quad (\text{A.8})$$

Expanding, we obtain

$$h_{\text{corr,DTS}}(f; g) = \frac{1}{1 + g \cdot \frac{\exp(-2j\pi f(T + \tau))}{1 - \exp(-2j\pi fT)}}, \quad (\text{A.9})$$

that is, the shorthand $h(f; g)$ of Eq. (3).

With the closed-loop transfer functions with a CTS and a DTS representation of the AO, we now study the critical behavior of the system as the gain increases to approach divergence.

A.2. Critical point. Analog case

The critical frequency and gain (Sect. 3.1) are determined by finding the couple of g_{crit} and f_{crit} that zero the denominator of Eq. (A.7). The equation to be solved for g and f is

$$g \frac{1}{2j\pi fT} \exp(-j\pi fT) \frac{\sin(\pi fT)}{\pi fT} \exp(-2j\pi f\tau) = -1. \quad (\text{A.10})$$

Equating for phase and modulus,

$$g \frac{\sin(\pi fT)}{2\pi^2 f^2 T^2} = 1, \quad \text{and} \quad -\frac{\pi}{2} - \pi fT - 2\pi f\tau = -\pi, \quad (\text{A.11})$$

which reduce to

$$f_{\text{crit}} = \frac{1}{2T + 4\tau} \quad (\text{A.12})$$

$$g_{\text{crit}} = \frac{2\pi^2 f_{\text{crit}}^2 T^2}{\sin(\pi f_{\text{crit}}T)}. \quad (\text{A.13})$$

A.3. Critical point. Digital case

In the same way as for Appendix A.2, the equation to be solved is

$$g \frac{\exp(-2j\pi f(T + \tau))}{1 - \exp(-2j\pi fT)} = -1. \quad (\text{A.14})$$

Leveraging

$$\begin{aligned} 1 - \exp(-2j\pi fT) &= 2j \exp(-j\pi fT) \sin(\pi fT) \\ &= 2 \exp\left(j\frac{\pi}{2}\right) \exp(-j\pi fT) \sin(\pi fT), \end{aligned}$$

we obtain the phase and modulus equations,

$$g = 2 \sin(\pi fT), \text{ and } -2\pi f\tau - 2\pi fT = -\pi + \pi/2 - \pi fT, \quad (\text{A.15})$$

and ultimately,

$$f_{\text{crit}} = \frac{1}{2T + 4\tau} \quad (\text{A.16})$$

$$g_{\text{crit}} = 2 \sin(\pi f_{\text{crit}}T), \quad (\text{A.17})$$

which is the general formula given in Table 1. We finally note that for a CTS with latency $\tau \geq 1$, we have $f_{\text{crit}}T < 1/6 \ll 1$, and therefore Eqs. (A.13) and (A.17) both reduce to the same value $g_{\text{crit}} \approx 2\pi f_{\text{crit}}T$ in a first-order approximation. The general discussions and derivations presented in Sect. 3 are in this case also applicable to a continuous-time description of the system. As a final note, Kulcsár et al. (2006) described the necessary hypotheses and proved that a discrete-time treatment of the minimization criterion in Eq. (2) is equivalent to a continuous-time treatment.

Appendix B: Semianalytical CLOSE solutions for latencies of 0 and 1 frames

We propose in this appendix some additional data to expand the discussion in Sect. 3.3 by exploring the close match obtained for all sensitivity-corrected $S/N \sigma_i$ between the minimum variance gain g_{MV} and the steady-state gain yielded by CLOSE for a setpoint $r = 0$. While in Fig. 6 this comparison was restricted to a latency $\tau = 2T$, we show in Fig. B.1 the same analysis for latencies $\tau = 0$ and $\tau = T$. All parameters are otherwise similar, in particular, the sample atmospheric spectrum (Eq. (1)) with a cutoff of 1 Hz. The time difference δk_{crit} at which the autocorrelation is locked to the setpoint r is taken from Table 1.

For all three latencies $\tau = 0, T$, and $2T$, we observe the same remarkable match between the minimum variance solution and CLOSE for a null setpoint. This was also simulated for integer latencies up to $\tau = 5T$, although these cases are not necessarily relevant for a real AO system and are not reported here. While Fig. B.1 (top) reports a very satisfactory correspondence, the analysis we proposed reaches the limit of the discrete time approximation of real AO systems, as we showed in Appendix A. With $f_{\text{crit}} = f_s/2$, the case $\tau = 0$ is therefore only adequate for simulated systems.

As commented on in the main text, the same conclusions apply at other latencies: Positive setpoints induce an undersetting of the integrator gain, while negative setpoints produce an over-setting and impose a gain floor at low S/N . This concludes our analysis of the minimum-variance capability of CLOSE for various latencies. We are confident that this property would extend in between, to fractional latencies within that range and beyond.

Fig. B.1. Minimum variance gain g_{MV} and solutions found using CLOSE with five different setpoint values, depending on the sensitivity corrected $S/N \sigma_i$. These computations are performed using the modal input spectrum shown in Fig. 3.

Appendix C: Experimenting with the correlation-locking setpoint – additional data

In this appendix we propose supplementary simulations to those in Sect. 5.3 to investigate the final long-exposure SR that is produced in the system imager versus the setpoint r used to drive CLOSE. The experimental protocol is identical to that in Sect. 5.3, generalized to three different latency values $\tau/T = 0, 1$, and 2 frames for three different AO system designs: the MICADO parameters used in the main text, and a PWFS and a Shack–Hartmann, both on an 8 m diameter telescope. The parameters for these simulations, thereafter referred to as PWFS8M and SH8M, are given in Table C.1 or are identical to those in Table 2 if unspecified.

Results with varying latencies, guide-star magnitudes, seeing conditions, and wind speeds against the r setpoint are given in Tables C.2–C.4, which also include results presented in Sect. 5.3 (see Table C.2, “Latency $\tau = 2T$ ”). In the tables presented here, we provide the maximum SR in H band for each given set of conditions [AO setup, τ , r_0 , M_R , wind speed V] that was achieved over r values from -0.5 to 0.5 . This value is noted $\text{SR}(r_{\text{max}})$, and the argument at which it is attained is noted r_{max} . For $r_{\text{max}} \neq 0$, we provide the value that was attained for the same parameters at our favorite all-in-one solution $r = 0$. This value is noted $\text{SR}(r = 0)$.

Conclusions for the MICADO simulations (Table C.2) are essentially identical to the analysis provided in Sect. 5.3: minor

fluctuations in the $r_{\max}$ value, with $r = -0.1$ or $r = 0$ being almost always optimal within the LE SR error bar of a few percent. These simulations are therefore conclusive in showing the versatility of CLOSE for different system latencies without additional complexities other than changing the value of δk_{crit} .

The simulations performed with the PWFS8M setup (Table C.2) show very similar trends to the MICADO setup, which indicates that our conclusion that $r = 0$ almost always optimally works may very well apply to a wide range of PWFS designs. A performance improvement and the automation of the integrator tune-up can also be a benefit for AO systems smaller than an ELT. The PWFS8M simulations at $\tau = 0$, $M_R = 12$ allow us to note the broadening of the bell function of the SR as a function of r , as was visible in Fig. 10. Some $r_{\max}$ values are shifted toward positive values, up to +0.3, but the difference between $\text{SR}(r_{\max})$ and $\text{SR}(r = 0)$ is negligible. It is lower than 1%. This behavior is expected because the margin for setting $\alpha_i G_i$, thus r , decreases with increased latency, as the trade-off between low frequency rejection and noise amplification becomes tighter. Thus, for very short latencies, the precision of $\alpha_i G_i$ for stable control and variance minimization is of little importance, and it follows that the setpoint r also benefits from an increased margin.

The simulations with the SH8M setup reported in Table C.4 are comparable to PWFSs, except for the brighter cases $M_R = 8$. For these, and most particularly for $M_R = 8$, $\tau = 0$, the ideal setpoint $r_{\max}$ is noticeably pushed to positive values, even hitting the extremum of our parameter space $r = +0.5$, although based on the recorded SRs, we may safely assume that no significant improvement would be reached for r in the 0.5–1.0 range. However, the difference in performance between $r_{\max} \geq 0.5$ and $r = 0$ is again insignificant, except for the cases with a wind speed of 40 m s^{-1} .

The cases discussed here (SH8M, $\tau = 0$, $M_R = 8$) are the one simulation in which modal optimization is least important, if at all: the sensor is linear, the S/N is extremely good, and the bandwidth effect is minimum. The need for balance between noise amplification and turbulence rejection is quite different from the design case explored in the main text. For all other cases with the SH8M setup, that is, where variance minimization by the rejection or noise trade-off is valid, we validated the near-minimum variance performance of CLOSE with a null setpoint, extending the conclusions seen for PWFSs. In particular, this also implies

Table C.1. AO numerical simulation parameters for simulation setups PWFS8M and SH8M.

Numerical simulation configurations		
	PWFS8M	SH8M
Telescope	$D = 8 \text{ m diameter}$ Circular pupil (no obstruction, no spider)	
DMs	Tip-tilt mirror Cartesian grid DM defined per: Pitch of 50 cm, coupling of 0.24 221 controlled actuators	
WFS		
Wavelength	658 nm	500 nm
Throughput	0.28	0.50
Subapertures	24×24	16×16 ^(a)
Measurements	1936 pixels total	368 slopes total
Modulation	$4 \frac{\lambda}{D}$ radius	–
Centroiding	–	Center of gravity
	–	300 mas px^{-1}
Readout noise	$0.3 e^-$	$3.0 e^-$ ^(b)
RTC controller		
Modes controlled	220	200 ^(a)
Frequency		$f_s = 500 \text{ Hz}$
Latency		$\tau = [0, T, 2T]$
CLOSE	Real-time implementation $p = 0.3$; $q^\pm = 10^{-3}$; $\delta k_{\text{crit}} = [1, 3, 5]$ $r = 0 [-0.5, 0.5]$	

Notes. Unspecified parameters are identical to the MICADO setup detailed in Table 2. ^(a)Fried geometry, with 20 highest-order KL modes filtered. ^(b)The read-out noise penalty for SH is loosely based on generally available detectors given the required number of pixels. This paper does not intend to compare Shack–Hartmann and PWFS. Guide-star magnitudes for the SH simulations are adjusted for a comparable S/N.

that the method we proposed is sound beyond certain hypotheses of the numerical analysis of Sect. 3, and in particular beyond the spectrum proposed in Eq. (1). The power law in Eq. (1) ought to be changed to $-11/3$ when accounting for the aliasing with a Shack–Hartmann sensor (Conan et al. 1995).

Table C.2. Summary of long-exposure SRs in H band for stationary simulations exploring the latency, seeing conditions, guide-star magnitude, and wind speed against the CLOSE r setpoint for the MICADO simulation setup.

Simulation setup: MICADO													
Latency $\tau = 0$													
r_0	M_R	12				15				16			
(cm)	V (m s ⁻¹)	10	20	30	40	10	20	30	40	10	20	30	40
16.3	SR($r_{\max}$)	84.0	83.2	82.7	81.6	75.5	69.4	63.8	58.4	62.4	46.6	36.3	30.7
	$r_{\max}$	-0.1	0.0	0.0	+0.1	0.0	-0.1	-0.1	0.0	0.0	-0.1	0.0	-0.1
	SR($r = 0$)	83.9			81.4		68.7	62.8			45.6		28.7
12.8	SR($r_{\max}$)	74.8	73.3	72.3	70.3	61.6	52.0	44.7	38.3	43.1	25.3	17.4	13.8
	$r_{\max}$	-0.1	-0.1	0.0	+0.1	0.0	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1
	SR($r = 0$)	74.7	73.2		70.2		50.0	42.4	37.5	41.7	22.8	15.4	11.8
8.9	SR($r_{\max}$)	50.4	48.5	45.0	41.5	29.6	20.6	14.3	10.2	13.6	4.2	2.2	0.9
	$r_{\max}$	0.0	-0.1	-0.1	0.0	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1
	SR($r = 0$)		47.8	44.6		27.8	16.3	11.2	6.7	8.3	1.6	0.8	0.5
Latency $\tau = T$													
r_0	M_R	12				15				16			
(cm)	V (m s ⁻¹)	10	20	30	40	10	20	30	40	10	20	30	40
16.3	SR($r_{\max}$)	82.3	79.4	74.6	64.4	72.7	61.6	50.0	39.4	58.8	36.6	23.4	18.1
	$r_{\max}$	0.0	0.0	-0.1	-0.1	0.0	-0.1	-0.1	-0.1	0.0	-0.1	-0.1	-0.1
	SR($r = 0$)			74.1	64.1		60.4	47.8	37.8		35.9	21.6	16.9
12.8	SR($r_{\max}$)	72.0	66.1	58.7	47.8	57.7	41.7	30.1	21.6	36.9	19.2	10.2	6.4
	$r_{\max}$	-0.1	0.0	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1
	SR($r = 0$)	71.8		57.7	44.3	56.4	39.6	26.9	20.2	35.9	13.9	9.2	5.0
8.9	SR($r_{\max}$)	46.5	38.0	29.9	20.5	24.4	12.7	6.8	3.8	9.6	1.9	0.5	0.3
	$r_{\max}$	-0.2	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1
	SR($r = 0$)	45.9	37.4	26.8	19.0	21.1	10.2	4.9	2.2	7.9	0.8	0.4	0.2
Latency $\tau = 2T$													
r_0	M_R	12				15				16			
(cm)	V (m s ⁻¹)	10	20	30	40	10	20	30	40	10	20	30	40
16.3	SR($r_{\max}$)	80.4	68.7	48.8	37.5	69.7	49.3	31.3	22.6	53.8	28.2	15.7	10.2
	$r_{\max}$	0.0	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	0.0	0.0	-0.1	-0.1	0.0
	SR($r = 0$)		67.2	44.3	36.5	49.1	46.2	30.9			27.0	15.5	
12.8	SR($r_{\max}$)	68.4	51.1	31.3	22.5	52.5	31.5	16.8	9.7	30.9	12.5	5.5	3.2
	$r_{\max}$	-0.1	-0.1	-0.2	-0.1	-0.1	-0.1	-0.1	0.0	-0.1	-0.1	-0.1	-0.1
	SR($r = 0$)	67.9	47.6	28.6	21.5	51.1	25.2	14.7		30.6	9.6	4.2	2.1
8.9	SR($r_{\max}$)	40.4	24.6	11.2	6.0	19.9	7.3	2.9	1.2	6.6	0.9	0.3	0.1
	$r_{\max}$	-0.1	-0.1	-0.2	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	0.0
	SR($r = 0$)	39.4	20.1	10.7	5.9	18.5	4.9	1.5	0.8	4.9	0.4	0.2	

Notes. $r_{\max}$: setpoint at which the maximum SR is obtained for a given parameter set. SR($r_{\max}$): maximum SR reached. SR($r = 0$): SR reached for a null setpoint when $r_{\max} \neq 0$.

Table C.3. Identical to Table C.2 for the PWFS8M simulation setup as defined in Table C.1.

Simulation setup: PWFS8M													
Latency $\tau = 0$													
r_0	M_R	12				15				16			
(cm)	$V \text{ (m s}^{-1}\text{)}$	10	20	30	40	10	20	30	40	10	20	30	40
16.3	$SR(r_{\max})$	83.4	83.3	82.8	81.9	76.9	73.7	70.4	67.0	68.4	61.2	54.1	48.4
	$r_{\max}$	0.0	+0.1	+0.2	+0.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	$SR(r = 0)$		83.2	82.4	81.0								
12.8	$SR(r_{\max})$	75.0	74.9	73.9	72.7	66.0	61.3	56.8	52.8	54.0	44.8	38.0	31.4
	$r_{\max}$	-0.1	+0.1	+0.2	+0.2	-0.1	0.0	-0.1	0.0	0.0	-0.1	-0.1	-0.1
	$SR(r = 0)$	74.9	74.8	73.4	71.7	65.3		56.7			44.4	36.5	30.7
8.9	$SR(r_{\max})$	55.3	55.0	52.9	51.5	41.7	36.6	30.3	25.9	28.0	19.5	13.5	9.2
	$r_{\max}$	-0.1	+0.1	+0.1	+0.2	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1	-0.1
	$SR(r = 0)$	55.1	54.9	52.7	50.8	40.6	35.3	29.2	24.8	25.8	17.5	11.8	8.4
Latency $\tau = T$													
r_0	M_R	12				15				16			
(cm)	$V \text{ (m s}^{-1}\text{)}$	10	20	30	40	10	20	30	40	10	20	30	40
16.3	$SR(r_{\max})$	82.4	81.1	78.3	73.3	74.9	69.1	62.4	54.4	66.1	55.0	45.7	36.2
	$r_{\max}$	0.0	+0.1	+0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	$SR(r = 0)$		81.0	78.1									
12.8	$SR(r_{\max})$	73.5	71.4	67.0	59.9	63.1	55.5	46.9	37.7	51.2	38.5	29.0	20.4
	$r_{\max}$	0.0	0.1	0.1	0.0	-0.1	0.0	0.0	0.0	0.0	-0.1	0.0	-0.1
	$SR(r = 0)$		71.3	66.8		63.0					38.3		20.3
8.9	$SR(r_{\max})$	52.9	49.3	42.3	33.7	38.9	29.2	20.3	13.2	24.6	14.2	7.5	4.2
	$r_{\max}$	0.0	0.0	0.0	0.0	-0.1	-0.1	-0.1	0.0	0.0	-0.1	-0.1	-0.1
	$SR(r = 0)$					37.9	28.2	19.8			12.9	7.1	4.1
Latency $\tau = 2T$													
r_0	M_R	12				15				16			
(cm)	$V \text{ (m s}^{-1}\text{)}$	10	20	30	40	10	20	30	40	10	20	30	40
16.3	$SR(r_{\max})$	81.4	76.6	65.5	51.2	72.9	63.3	50.2	37.6	63.5	48.4	35.3	25.3
	$r_{\max}$	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	$SR(r = 0)$												
12.8	$SR(r_{\max})$	71.7	64.2	50.1	34.8	60.7	47.5	33.6	21.5	48.0	31.2	19.6	12.0
	$r_{\max}$	0.0	0.0	0.0	0.0	0.0	0.0	-0.1	0.0	0.0	0.0	0.0	0.0
	$SR(r = 0)$							33.2					
8.9	$SR(r_{\max})$	49.7	38.9	24.0	12.7	35.1	21.5	10.9	5.3	20.9	9.1	4.1	1.9
	$r_{\max}$	0.0	0.0	-0.1	0.0	-0.1	-0.1	-0.1	0.0	-0.1	0.0	0.0	0.0
	$SR(r = 0)$			23.9		34.7	21.4	10.7		20.7			

Table C.4. Identical to Table C.2 for the SH8M simulation setup as defined in Table C.1.

Simulation setup: SH8M													
Latency $\tau = 0$													
r_0	M_R	8				11				12			
(cm)	$V \text{ (m s}^{-1}\text{)}$	10	20	30	40	10	20	30	40	10	20	30	40
16.3	$\text{SR}(r_{\max})$	80.8	80.8	81.1	80.9	76.3	74.0	72.3	71.1	63.7	56.3	50.9	47.7
	$r_{\max}$	+0.3	+0.5*	+0.5*	+0.5*	0.0	0.0	0.0	+0.1	0.0	0.0	0.0	0.0
	$\text{SR}(r = 0)$	80.7	80.3	79.9	78.7				70.8				
12.8	$\text{SR}(r_{\max})$	73.0	73.1	73.4	73.2	68.1	65.6	63.8	62.7	55.1	47.0	43.0	39.5
	$r_{\max}$	+0.4	+0.5*	+0.5*	+0.5*	0.0	0.0	0.0	+0.1	0.0	0.0	0.0	0.0
	$\text{SR}(r = 0)$	72.9	72.3	71.8	70.3				62.1				
8.9	$\text{SR}(r_{\max})$	57.2	57.4	57.8	57.5	52.0	49.6	48.7	47.4	40.4	32.9	29.1	26.2
	$r_{\max}$	+0.3	+0.5*	+0.5*	+0.5*	0.0	0.0	+0.1	+0.1	0.0	0.0	0.0	0.0
	$\text{SR}(r = 0)$	57.0	56.4	55.5	53.4			48.1	46.4				
Latency $\tau = T$													
r_0	M_R	8				11				12			
(cm)	$V \text{ (m s}^{-1}\text{)}$	10	20	30	40	10	20	30	40	10	20	30	40
16.3	$\text{SR}(r_{\max})$	80.2	79.1	77.4	74.0	74.9	70.4	66.4	62.8	61.6	51.9	44.4	39.3
	$r_{\max}$	+0.3	+0.3	+0.2	+0.1	0.0	0.0	0.0	0.0	+0.1	0.0	0.0	0.0
	$\text{SR}(r = 0)$	80.1	78.6	76.9	73.9					61.5			
12.8	$\text{SR}(r_{\max})$	72.2	70.9	68.7	64.3	66.2	61.3	57.3	52.9	52.4	42.6	35.8	30.3
	$r_{\max}$	+0.3	+0.3	+0.2	+0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	$\text{SR}(r = 0)$	72.0	70.2	67.9	64.1								
8.9	$\text{SR}(r_{\max})$	56.1	54.3	51.4	45.6	50.1	45.1	40.9	35.9	37.5	28.2	22.5	18.2
	$r_{\max}$	+0.3	+0.3	+0.2	+0.1	0.0	+0.1	+0.1	0.0	0.0	0.0	0.0	0.0
	$\text{SR}(r = 0)$	55.8	53.4	50.4	45.3		45.0	40.7					
Latency $\tau = 2T$													
r_0	M_R	8				11				12			
(cm)	$V \text{ (m s}^{-1}\text{)}$	10	20	30	40	10	20	30	40	10	20	30	40
16.3	$\text{SR}(r_{\max})$	79.2	74.6	66.3	58.0	73.1	65.8	57.0	49.0	59.4	46.6	37.3	30.2
	$r_{\max}$	+0.2	+0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	$\text{SR}(r = 0)$	79.0	74.5										
12.8	$\text{SR}(r_{\max})$	70.9	65.0	54.4	44.5	64.2	55.5	46.2	37.4	49.8	36.3	28.0	21.1
	$r_{\max}$	+0.2	+0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	$\text{SR}(r = 0)$	70.5	74.9										
8.9	$\text{SR}(r_{\max})$	54.3	46.7	33.8	23.6	47.6	38.4	27.6	19.2	34.6	22.7	14.8	9.6
	$r_{\max}$	+0.2	+0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	$\text{SR}(r = 0)$	53.7	46.4										

Notes. Values with an asterisk “ $r_{\max} = +0.5^*$ ” indicate that the actual maximum was beyond our probing range, and thus only guarantees that $r_{\max} \geq 0.5$. In these cases, the value reported as $\text{SR}(r_{\max})$ was achieved for $r = +0.5$ and is not the actual maximum.