First on-sky SCAO validation of full LQG control with vibration mitigation on the CANARY pathfinder

Gaetano Sivo, Caroline Kulcsár, Jean-Marc Conan, Henri-François Raynaud, Eric Gendron, Alastair Basden, Fabrice Vidal, Tim Morris, Serge Meimon, Cyril Petit, et al.

HAL Id: hal-01739389

https://hal-iogs.archives-ouvertes.fr/hal-01739389

Submitted on 22 Dec 2023

First on-sky SCAO validation of full LQG control with vibration mitigation on the CANARY pathfinder

Gaetano Sivo,[1,2,7] Caroline Kulcsár,[1,*] Jean-Marc Conan,[2] Henri-François Raynaud,[1] Éric Gendron,[3] Alastair Basden,[4] Fabrice Vidal,[3] Tim Morris,[4] Serge Meimon,[2] Cyril Petit,[2] Damien Gratadour,[3] Olivier Martin,[3] Zoltan Hubert,[3] Arnaud Sevin,[3] Denis Perret,[3] Fanny Chemla,[5] Gérard Rousset,[3] Nigel Dipper,[4] Gordon Talbot,[4] Eddy Younger,[4] Richard Myers,[4] David Henry,[6] Stephen Todd,[6] David Atkinson,[6] Colin Dickson,[6] and Andy Longmore[6]

[1]*Laboratoire Charles Fabry, Institut d'Optique Graduate School-CNRS, Palaiseau, France*
[2]*ONERA, The French Aerospace Lab, Châtillon, France*
[3]*Laboratoire d'Étude Spatiale et d'Instrumentation en Astrophysique, Observatoire de Paris, Meudon, France*
[4]*Centre for Advanced Instrumentation, Durham University, South Road, Durham, UK*
[5]*Galaxies Étoiles Planètes et Instrumentation, Observatoire de Paris, Meudon, France*
[6]*UK Astronomy Technology Centre, Royal Observatory Edinburgh, Blackford Hill, Edinburgh EH9 3HJ, UK*
[7]*Université Paris 13, Villetaneuse, France*
[*]caroline.kulcsar@institutoptique.fr

Abstract: Adaptive optics provides real time correction of wavefront disturbances on ground based telescopes. Optimizing control and performance is a key issue for ever more demanding instruments on ever larger telescopes affected not only by atmospheric turbulence, but also by vibrations, windshake and tracking errors. Linear Quadratic Gaussian control achieves optimal correction when provided with a temporal model of the disturbance. We present in this paper the first on-sky results of a Kalman filter based LQG control with vibration mitigation on the CANARY instrument at the Nasmyth platform of the 4.2-m William Herschel Telescope. The results demonstrate a clear improvement of performance for full LQG compared with standard integrator control, and assess the additional improvement brought by vibration filtering with a tip-tilt model identified from on-sky data, thus validating the strategy retained on the instrument SPHERE at the VLT.

References and links

1. G. Rousset, J.-C. Fontanella, P. Y. Kern, P. Gigan, F. Rigaut, P. Léna, C. Boyer, P. Jagourel, J.-P. Gaffard, and F. Merkle, "First diffraction-limited astronomical images with adaptive optics," Astron. Astrophys. **230**, 29–32 (1990).
2. R. M. Myers, Z. Hubert, T. J. Morris, E. Gendron, N. A. Dipper, A. Kellerer, S. J. Goodsell, G. Rousset, E. Younger, M. Marteaud, A. G. Basden, F. Chemla, C. D. Guzman, T. Fusco, D. Geng, B. Le Roux, M. A. Harrison,

A. J. Longmore, L. K. Young, F. Vidal, and A. H. Greenaway, "CANARY: the on-sky NGS/LGS MOAO demonstrator for EAGLE," Proc. SPIE **7015**, 70150E (2008).

3. F. Hammer, F. Sayède, E. Gendron, T. Fusco, D. Burgarella, V. Cayatte, J.-M. Conan, F. Courbin, H. Flores, I. Guinouard, L. Jocou, A. Lançon, G. Monnet, M. Mouhcine, F. Rigaud, D. Rouan, G. Rousset, V. Buat, and F. Zamkotsian, "The FALCON concept: multi-object spectroscopy combined with MCAO in near-IR," in *Scientific Drivers for ESO Future VLT/VLTI Instrumentation*, ESO Astrophysics Symposia, J. Bergeron and G. Monnet, eds. (Springer, 2002), pp. 139–148.

4. J.-G. Cuby, S. Morris, T. Fusco, M. Lehnert, P. Parr-Burman, G. Rousset, J.-P. Amans, S. Beard, I. Bryson, M. Cohen, N. Dipper, C. Evans, M. Ferrari, E. Gendron, J.-L. Gimenez, D. Gratadour, P. Hastings, Z. Hubert, E. Hugot, P. Jagourel, P. Laporte, V. Lebrun, D. L. Mignant, F. Madec, R. Myers, B. N. T. Morris, C. Robert, H. Schnetler, M. Swinbank, G. Talbot, W. Taylor, F. Vidal, S. Vivès, P. Vola, N. Welikala, and M. Wells, "EAGLE: a MOAO fed multi-IFU NIR workhorse for E-ELT," Proc. SPIE **7735**, 77352D (2010).

5. E. Gendron, F. Vidal, M. Brangier, T. Morris, Z. Hubert, A. Basden, G. Rousset, R. Myers, F. Chemla, A. Longmore, T. Butterley, N. Dipper, C. Dunlop, D. Geng, D. Gratadour, D. Henry, P. Laporte, N. Looker, D. Perret, A. Sevin, G. Talbot, and E. Younger, "MOAO first on-sky demonstration with CANARY," Astron. Astrophys. **529**, L2 (2011).

6. C. Kulcsár, H.-F. Raynaud, C. Petit, J.-M. Conan, and P. Viaris de Lesegno, "Optimal control, observers and integrators in adaptive optics," Opt. Express **14**, 7463–8012 (2006).

7. C. Kulcsár, H.-F. Raynaud, C. Petit, and J.-M. Conan, "Minimum variance prediction and control for adaptive optics," Automatica **48**, 1939–1954 (2012).

8. B. Le Roux, J.-M. Conan, C. Kulcsár, H.-F. Raynaud, L. M. Mugnier, and T. Fusco, "Optimal control law for classical and multiconjugate adaptive optics," J. Opt. Soc. Am. A **21**, 1261–1276 (2004).

9. C. Petit, J.-M. Conan, C. Kulcsár, and H.-F. Raynaud, "Linear quadratic gaussian control for adaptive optics and multiconjugate adaptive optics: experimental and numerical analysis," J. Opt. Soc. Am. A **26**, 1307–1325 (2009).

10. C. Kulcsár, H.-F. Raynaud, C. Petit, and J.-M. Conan, "Minimum variance control in presence of actuator saturation in adaptive optics," Proc. SPIE **7015**, 70151G (2008).

11. J.-M. Conan, H.-F. Raynaud, C. Kulcsár, and S. Meimon, "Are integral controllers adapted to the new era of ELT adaptive optics?" in *2nd International Conference on Adaptive Optics for Extremely Large Telescopes (AO4ELT)* (Victoria, Canada, 2011).

12. C. Kulcsár, G. Sivo, H.-F. Raynaud, B. Neichel, F. Rigaut, J. Christou, A. Guesalaga, C. Correia, J.-P. Véran, E. Gendron, F. Vidal, G. Rousset, T. Morris, S. Esposito, F. Quirós-Pacheco, G. Agapito, E. Fedrigo, L. Pettazzi, R. Clare, R. Muradore, S. Meimon, and J.-M. Conan, "Vibrations in AO control: a short analysis of on-sky data around the world," Proc. SPIE **8447**, 84471C (2012).

13. C. Correia, J.-P. Véran, and G. Herriot, "Advanced vibration suppression algorithms in adaptive optics systems," J. Opt. Soc. Am. A **29**, 185–194 (2012).

14. G. Agapito, G. Battistelli, D. Mari, D. Selvi, A. Tesi, and P. Tesi, "Frequency based design of modal controllers for adaptive optics systems," Opt. Express **20**, 27108–27122 (2012).

15. A. Guesalaga, B. Neichel, J. O'Neal, and D. Guzman, "Mitigation of vibrations in adaptive optics by minimization of closed-loop residuals," Opt. Express **21**, 10676–10696 (2013).

16. C. Petit, J.-M. Conan, C. Kulcsár, H.-F. Raynaud, and T. Fusco, "First laboratory validation of vibration filtering with LQG control law for adaptive optics," Opt. Express **16**, 87–97 (2008).

17. S. Meimon, C. Petit, T. Fusco, and C. Kulcsár, "Tip-tilt disturbance model identification for Kalman-based control scheme: application to XAO and ELT systems," J. Opt. Soc. Am. A **27**, A122–A132 (2010).

18. C. Petit, T. Fusco, J.-F. Sauvage, A. Sevin, M. Suarez, J. Charton, P. Baudoz, J.-L. Beuzit, T. Buey, K. Dohlen, P. Feautrier, E. Fedrigo, B. Fleury, J.-L. Gach, N. Hubin, M. Kasper, D. Mouillet, D. Perret, P. Puget, J.-C. Sinquin, C. Soenke, and F. Wildi, "Integration, tests and laboratory performance of SAXO, the VLT-SPHERE extreme AO system," in *2nd International Conference on Adaptive Optics for Extremely Large Telescopes (AO4ELT)* (2011).

19. L. A. Poyneer and J.-P. Véran, "Kalman filtering to suppress spurious signals in adaptive optics control," J. Opt. Soc. Am. A **27**, A223–A234 (2010).

20. A. Costille, C. Petit, J.-M. Conan, C. Kulcsár, H.-F. Raynaud, and T. Fusco, "Wide field adaptive optics laboratory demonstration with closed-loop tomographic control," J. Opt. Soc. Am. A **27**, 469–483 (2010).

21. N. Doelman, R. Fraanje, and R. Breeje, "Real-sky adaptive optics experiments on optimal control of tip-tilt modes," in *2nd International Conference on Adaptive Optics for Extremely Large Telescopes (AO4ELT)* (2011).

22. G. Sivo, C. Kulcsár, J.-M. Conan, H.-F. Raynaud, E. Gendron, A. Basden, F. Vidal, T. Morris, S. Meimon, C. Petit, D. Gratadour, O. Martin, Z. Hubert, G. Rousset, N. Dipper, G. Talbot, E. Younger, and R. Myers, "Full LQG control with vibration mitigation: From theory to first on-sky validation on the canary MOAO demonstrator," in *Imaging and Applied Optics* (Optical Society of America, 2013), p. OTu2A.2.

23. G. Sivo, C. Kulcsár, J.-M. Conan, H.-F. Raynaud, E. Gendron, A. Basden, F. Vidal, T. Morris, S. Meimon, C. Petit, D. Gratadour, O. Martin, Z. Hubert, G. Rousset, N. Dipper, G. Talbot, E. Younger, and R. Myers, "First on-sky validation of full LQG control with vibration mitigation on the CANARY MOAO pathfinder," in *3rd International Conference on Adaptive Optics for Extremely Large Telescopes (AO4ELT)* (Florence, Italy, 2013).

24. T. Morris, Z. Hubert, R. Myers, E. Gendron, A. Longmore, G. Rousset, G. Talbot, T. Fusco, N. Dipper, F. Vidal, D. Henry, D. Gratadour, T. Butterley, F. Chemla, D. Guzman, P. Laporte, E. Younger, A. Kellerer, M. Harrison, M. Marteaud, D. Geng, A. Basden, A. Guesalaga, C. Dunlop, S. Todd, C. Robert, K. Dee, C. Dickson, N. Vedrenne, A. Greenaway, B. Stobie, H. Dalgarno, and J. Skvarc, "CANARY: The NGS/LGS MOAO demonstrator for EAGLE," in *1st International Conference on Adaptive Optics for Extremely Large Telescopes (AO4ELT)* (2009).

25. E. Gendron, T. Morris, Z. Hubert, R. Myers, A. Longmore, G. Rousset, G. Talbot, F. Vidal, N. Dipper, D. Gratadour, N. Looker, M. Brangier, E. Younger, A. Sevin, A. Basden, D. Perret, L. Young, D. Atkinson, F. Chemla, D. Henry, T. Butterley, P. Laporte, D. Guzman, M. Marteaud, D. Geng, N. Vedrenne, M. Harrison, T. Fusco, A. Guesalaga, C. Dunlop, S. Todd, K. Dee, C. Dickson, A. Greenaway, B. Stobie, H. Dalgarno, and J. Skvarc, "Status update of the CANARY on-sky MOAO demonstrator," Proc. SPIE **7736**, 77360P (2010).

26. A. Basden, D. Geng, R. Myers, and E. Younger, "Durham adaptive optics real-time controller," Appl. Opt. **49**, 6354–6363 (2010).

27. A. G. Basden and R. M. Myers, "The Durham adaptive optics real-time controller: capability and Extremely Large Telescope suitability," Mon. Not. R. Astron. Soc. **424**, 1483–1494 (2012).

28. B. D. O. Anderson and J. B. Moore, *Optimal Control, Linear Quadratic Methods* (Prentice Hall, 1990).

29. B. Neichel, A. Parisot, C. Petit, T. Fusco, and F. Rigaut, "Identification and calibration of the interaction matrix parameters for AO and MCAO systems," Proc. SPIE **8447**, 84475N (2012).

30. A. Parisot, "Systèmes d'optique adaptative grand champ pour les ELT: validation expérimentale de nouveaux concepts sur le banc HOMER," Ph.D. thesis, Université de Provence Aix-Marseille 1 (2012).

31. F. Vidal, E. Gendron, and G. Rousset, "Tomography approach for multi-object adaptive optics," J. Opt. Soc. Am. A **27**, A253–A264 (2010).

32. C. Kulcsár, H.-F. Raynaud, J.-M. Conan, C. Correia, and C. Petit, "Control design and turbulent phase models in adaptive optics: a state-space interpretation," in *Adaptive Optics: Analysis and Methods/Signal Recovery and Synthesis Topical Meetings,* Optical Society of America, San José, CA, USA (2009).

33. G. Sivo, H.-F. Raynaud, J.-M. Conan, C. Kulcsár, E. Gendron, F. Vidal, and A. Basden, "First laboratory validation of LQG control with the CANARY MOAO pathfinder," Proc. SPIE **8447**, 84472Y (2012).

34. V. Kučera, *Analysis and Design of Discrete Linear Control Systems*, Series in Systems and Control Engineering (Prentice Hall, 1991).

35. C. Petit, "Development of a new control law for the adaptive optics facility (AOF)," Tech. Rep. RF 1/17838 DOTA, Onera (2011).

36. J.-M. Conan, G. Rousset, and P.-Y. Madec, "Wave-front temporal spectra in high-resolution imaging through turbulence," J. Opt. Soc. Am. A **12**, 1559–1570 (1995).

37. O. Martin, É. Gendron, G. Rousset, and F. Vidal, "Temporal convergence of phase spatial covariance matrix measurements in tomographic adaptive optics," Proc. SPIE **8447**, 84472A (2012).

38. C. Béchet, M. Le Louarn, M. Tallon, and E. Thiébaut, "Performances of the fractal iterative method with an internal model control law on the ESO end-to-end ELT adaptive optics simulator," Proc. SPIE **7015**, 70151H (2008).

39. C. Correia, H.-F. Raynaud, C. Kulcsár, and J.-M. Conan, "Minimum variance control for the woofer-tweeter concept," in *Adaptive Optics: Methods, Analysis and Applications/Signal Recovery and Synthesis Topical Meetings,* Optical Society of America, San José, CA (2009), p. AOWB4.

40. P. Massioni, C. Kulcsár, H.-F. Raynaud, and J.-M. Conan, "Fast computation of an optimal controller for large-scale adaptive optics," J. Opt. Soc. Am. A **28**, 2298–2309 (2011).

41. L. Gilles, P. Massioni, C. Kulcsár, H.-F. Raynaud, and B. Ellerbroek, "Distributed Kalman filtering compared to Fourier domain preconditioned conjugate gradient for laser guide star tomography on extremely large telescopes," J. Opt. Soc. Am. A **30**, 898–909 (2013).

42. M. M. Seron, J. H. Braslavsky, and G. C. Goodwin, *Fundamental Limitations in Filtering and Control* (Springer, 1997).

43. J. Osborn, R. Wilson, H. Shepherd, T. Butterley, V. Dihllon, and R. Avila, "Stereo scidar: Profiling atmospheric optical turbulence with improved altitude resolution," in *3rd International Conference on Adaptive Optics for Extremely Large Telescopes (AO4ELT)* (2013).

44. A. Cortés, B. Neichel, A. Guesalaga, J. Osborn, F. Rigaut, and D. Guzman, "Atmospheric turbulence profiling using multiple laser star wavefront sensors," Mon. Not. R. Astron. Soc. **427**, 2089–2099 (2012).

45. L. Poyneer, M. van Dam, and J.-P. Véran, "Experimental verification of the frozen flow atmospheric turbulence assumption with use of astronomical adaptive optics telemetry," J. Opt. Soc. Am. A **26**, 833–846 (2009).

46. L. Wang, M. Schöck, and G. Chanan, "Atmospheric turbulence profiling with slodar using multiple adaptive optics wavefront sensors," Appl. Opt. **47**, 1880–1892 (2008).

47. S. M. Ammons, L. Poyneer, D. T. Gavel, R. Kupke, C. E. Max, and L. Johnson, "Evidence that wind prediction with multiple guide stars reduces tomographic errors and expands moao field of regard," Proc. SPIE **8447**, 84471U (2012).

1. Introduction

Adaptive optics (AO) is a technique that allows a deformable mirror (DM) to perform real-time compensation of the wavefront (WF) distortion induced by the atmospheric turbulence encountered during astronomical observations. AO has been successfully demonstrated and made available for astronomers since the early 1990's [1]. Its success is such that the next generation of Extremely Large Telescopes (ELTs) cannot be conceived without it. The associated future astrophysical programs rely on sophisticated AO systems able to fulfill ever demanding requirements. These systems can be globally split into two categories. On one hand wide-field AO (WFAO) uses a large field for wavefront sensing and/or for correction. On the other hand single conjugated AO (SCAO) uses an on-axis guide star both for wavefront sensing and correction. Extreme AO corresponds to very high performance SCAO feeding specific high contrast instruments.

The present paper is focused on SCAO, with on-sky experiments conducted using the SCAO mode of the CANARY demonstrator [2]. This demonstrator has been initially designed to validate a specific WFAO concept: multi-object AO (MOAO), required for the study of primordial galaxies, where several objects of small angular dimensions and spread over a large field of view are observed simultaneously [3]. The CANARY MOAO pathfinder is the demonstrator for the future MOAO instrument EAGLE [4] and the first on-sky demonstration of MOAO was performed successfully in September 2010 [5].

Control and performance optimization of complex AO systems are a key issue. Linear Quadratic Gaussian (LQG) control is an appealing strategy for AO since it provides globally optimal control with respect to an explicit minimum variance performance criterion [6, 7]. Numerical studies [8] and laboratory demonstrations [9] have shown that LQG control can give better performance than integral controllers due to the use of spatial and temporal priors. Theoretical studies [10, 11] have also shown that an advantage of LQG controllers over integrator was the possible relaxation of system constraints (actuator saturation, sampling frequency, noise requirements) which may be of particular interest for ELT systems. A major issue in AO is also the correction of non-turbulent perturbations such as windshake or vibrations [12–15]. LQG control, thanks to its flexibility, can provide a very efficient compensation of these effects, as shown for vibrations through laboratory tests [16], and through numerical simulations for windshake [17].

LQG control has been implemented for tip-tilt correction on the SPHERE instrument, the European extrasolar planet finder [18]. It has also been considered during preliminary studies on GPI using a Kalman filter with predictive Fourier-domain control [19]. LQG control has been demonstrated on-lab for WFAO systems [20]. Note that tip tilt optimal control has been tested on-sky using an AO system installed at the McMath-Pierce solar telescope [21]. The next step was to demonstrate full LQG control on-sky, full meaning with optimal LQG control of all modes, and not only tip and tilt.

This paper presents a detailed analysis of the first SCAO on-sky validation of a full LQG control with vibration mitigation, performed in July and August 2012 [22, 23]. We compare performance obtained with integrator versus LQG, with and without vibration mitigation. Vibration models for tip and tilt modes are identified on line from WFS data, using the frequency-domain identification method selected for the instrument SPHERE [17], hence validating on sky and for the first time its tip-tilt identification and control strategy. Analysis of models relevance is conducted, and the advantage of LQG control is also discussed on the basis of power spectral densities and rejection transfer functions.

We organize the paper as follows. In Section 2 we briefly present the CANARY bench and its SCAO mode. Section 3 introduces the formalism of the LQG control for SCAO. In Section 4 we discuss practical implementation issues such as model identification or real-time complexity for

the CANARY instrument. Then in Section 5 we present the first on-sky results with a detailed data analysis. This is followed in Section 6 by a frequency-domain and a robustness analysis for integrator and LQG, and Section 7 concludes this paper.

2. The single conjugated adaptive optics mode of CANARY

CANARY is a European project led by Durham University and Observatoire de Paris. The CANARY bench is set up at the Nasmyth focus of the 4.2-m William Herschel Telescope at the summit of the Roque de los Muchachos mountain on La Palma, Canary Islands, Spain [24, 25]. CANARY can perform AO in several possible configurations including SCAO, MOAO using a four Natural Guide Stars (NGSs) asterism (three off-axis stars and one on-axis for performance evaluation and scientific IR images) and laser based MOAO with four additional Laser Guide Stars (LGSs) in the technical Field of View (FoV). In this paper we focus on the SCAO mode of CANARY. In SCAO, an on-axis WF Sensor (WFS) (so-called Truth Sensor (TS)) provides closed-loop 7×7 Shack Hartmann WFS data on the on-axis star. Correction is performed by a high-speed tip-tilt mirror developed by Observatoire de Paris, and a CILAS 8×8 piezostack array DM conjugated to the telescope's pupil. The CANARY Real-Time Computer (RTC), called DARC, is developed by Durham University [26, 27] and IR images are recorded thanks to the near-IR camera CAMICAz, developed by Observatoire de Paris. CAMICAz has a NICMOS-3 array composed of 256×256 pixels. The camera has several spectral filters. We observe in the two spectral bands H and K. The central wavelengths are $\lambda_H = 1.67 \mu$m and $\lambda_K = 2.18 \mu$m with $\Delta_{\lambda_H} = 0.14 \mu$m and $\Delta_{\lambda_K} = 0.15 \mu$m. The plate scale of the IR camera is $p = 28.8$ mas. A telescope simulator developed by UK ATC allows performance evaluation on internal sources with emulated turbulence being introduced by rotating phase plates.

3. LQG control for SCAO

LQG control can be derived when the system is described by a linear state space model. This approach provides a general framework based on explicit models of the AO system including wavefront perturbations [8]. We present here a full LQG on-sky SCAO demonstration, *i.e.* LQG control is applied on all modes (tip, tilt and higher orders). For tip and tilt modes, vibratory components are accounted for as in [16, 17].

In this Section we explain how to establish the dynamical models including vibrations, and we derive an optimal control based on a criterion that simplifies system calibrations. We first present the minimum variance performance criterion and the optimal solution in Section 3.1. We then present the modeling of the disturbance stochastic process (Section 3.2) with a specific treatment for tip-tilt components (Section 3.3). Section 3.4 establishes the corresponding SCAO state space representation on which is based the Kalman filter used for control computation (Section 3.5).

3.1. Performance criterion and optimal solution

The perturbation phase is a continuous time process. Since the DM is applied through a zero order hold and the DM's time response is small compared to the loop sampling period T, the continuous time minimum-variance criterion can be equivalently replaced by a discrete time criterion for control optimization [6]. We briefly recall here this construction in the simple 'synchronous' case where controls are applied at the same sampling times $t = kT$, and where one sampling interval is allocated for WFS measurement processing and control calculation. In this case, the control objective is therefore to minimize the following discrete-time infinite-horizon quadratic performance criterion $J(u)$, which corresponds to the variance of the residual

phase $\phi_k^{\text{res}} \triangleq \phi_k - \phi_k^{\text{cor}}$:

$$J(u) = \lim_{K \to +\infty} \frac{1}{K} \sum_{k=1}^{K} ||\phi_k^{\text{res}}||^2 = \lim_{K \to +\infty} \frac{1}{K} \sum_{k=1}^{K} ||\phi_{k+1} - Nu_k||^2 \tag{1}$$

where the correction phase ϕ_k^{cor} is related to the command vector u through

$$\phi_k^{\text{cor}} = Nu_{k-1}, \tag{2}$$

and where u_{k-1} is the control vector applied at time $(k-1)T$, N being the DM influence matrix. In formula (1), the discrete-time perturbation ϕ_k is the averaged perturbation over the sampling interval $[(k-1)T, kT[$, defined as

$$\phi_k \triangleq \frac{1}{T} \int_{(k-1)T}^{kT} \phi(t) \, dt. \tag{3}$$

All phase variables are defined similarly as temporal averages over one frame of duration T. When the phase is described by Zernike coefficients, the influence matrix N is the operator that connects command vector u to Zernike modes.

The optimal solution u_k^* that minimizes $J(u)$ is given through an LQG approach if these two hypothesis are satisfied:

- the system can be described by a linear state-space model with additional noise inputs;

- the noises are white and Gaussian, *i.e.* uncorrelated sequences of Gaussian vectors; co-variance matrices are known.

More details on LQG control can be found in [28]. When the control u_k is computed from measurements $y_k, \ldots, y_0$, the stochastic separation principle allows to solve this problem in two steps:

- compute the estimated phase as the conditional expectation of ϕ_{k+1} knowing all measurements $y_k, \ldots, y_0$ available at time step k:

$$\widehat{\phi}_{k+1|k} \triangleq \mathrm{E}[\phi_{k+1}|y_k, \ldots, y_0] \tag{4}$$

- project the estimated phase onto the DM:

$$u_k^* = P_u \widehat{\phi}_{k+1|k} \tag{5}$$

where P_u is the projector defined as pseudo-inverse of N:

$$P_u \triangleq N^{\dagger} = (N^{\mathsf{T}} N)^{-1} N^{\mathsf{T}} \tag{6}$$

The optimal prediction $\widehat{\phi}_{k+1|k}$ is computed recursively as the output of the Kalman filter based on a state-space representation of the complete system (see Section 3.5).

The measurement equation is defined as

$$y_k = D\phi_{k-1}^{\text{res}} + w_k \tag{7}$$

which can be written using Eq. (2) as

$$y_k = D\phi_{k-1} - DNu_{k-2} + w_k, \tag{8}$$

where D is the WFS matrix that characterizes the WFS response in terms of the input phase, and w is the measurement noise supposed to be white Gaussian and independent of ϕ, with covariance matrix Σ_w. Note that this measurement equation features a two-frame loop delay between DM control u and WFS output y, also called 'plant delay' in control literature. This loop delay induces an important limitation of the closed-loop performance.

A difficulty may arise for the calculation of the projector $N^\dagger$ in Eq. (6). It requires the knowledge of the influence matrix N, which depends on the actual system geometry. N is not only related to the DM influence functions but also to the relative position of the actuator grid with respect to the pupil. Moreover it has to be oriented with respect to the WFS geometry. Rotations, magnifications, and misregistrations may occur between these grids which imply specific calibration procedures [29, 30] that can also be avoided by modifying the control criterion $J(u)$ as explained below.

Since the interaction matrix defined as $M_{\text{int}} \triangleq DN$ is easily and regularly recorded on the bench, we propose to minimize the alternative quadratic criterion $\mathscr{L}$ defined as

$$\mathscr{L}(u) = \lim_{K \to +\infty} \frac{1}{K} \sum_{k=1}^{K} ||D\phi_{k+1} - DNu_k||^2 = \lim_{K \to +\infty} \frac{1}{K} \sum_{k=1}^{K} ||D\phi_k^{\text{res}}||^2. \tag{9}$$

This modified criterion is the residual slopes variance as opposed to the residual phase variance $J(u)$. Minimizing $\mathscr{L}(u)$ instead of $J(u)$ leads to the same optimal phase estimate and the same control structure, but with a different projector onto the DM, namely:

$$P_u \triangleq (DN)^\dagger D \tag{10}$$

Computing this new projector only requires the knowledge of the WFS matrix D and of the command matrix $M_{\text{com}} = (DN)^\dagger$, i.e. the pseudo-inverse of the command-to-slopes interaction matrix $M_{\text{int}} = DN$. Although sub-optimal with respect to the residual phase variance criterion $J(u)$, this control has the advantage to be very robust to system mis-alignments and avoids tricky calibration procedures for the determination of N. The relative geometry between WFS and DM is indeed accounted for in M_{int}, therefore avoiding an explicit determination of magnification, rotation and decentering between WFS subaperture grid and DM actuator grid. Note that the same 'slopes-oriented' approach lies at the heart of the APPLY MOAO control strategy, where open-loop on-axis WFS slopes are reconstructed from off-axis measurements using a static minimum-variance estimator [31].

The WFS matrix D is needed to compute the slopes-to-control projector P_u. Because it appears in the measurement equation (8), it is also needed to compute the Kalman filter presented in Section 3.5. In practice, D was computed theoretically from a geometrical model and scaled using straightforward calibrations done on internal sources, a procedure which enables to determine both the sensitivity of the WFS and the ordering of the subapertures. The interaction matrix $M_{\text{int}} = DN$ was also calibrated on the bench.

3.2. Disturbance stochastic dynamical model

As explained above, the computation of the optimal prediction $\widehat{\phi}_{k+1|k}$ relies on a state-space representation. It can be established when a stochastic dynamical model of the perturbation is available. Tip and tilt (TT) are the most energetic turbulent modes and frequently include additional perturbations due to vibrations in the telescope and system structure, windshake and residual tracking errors. These factors led us to separate TT modeling from that of the higher orders (HO).

Vibrations are additive perturbations and independent from the turbulence, so that the complete perturbation phase can be written as:

$$\phi = \phi^{\text{tur}} + \phi^{\text{vib}}. \tag{11}$$

Assuming that vibrations only affect the tip and tilt modes, the TT and HO components of the disturbance ϕ, written as

$$\phi = \begin{pmatrix} \phi^{\mathrm{TT}} \\ \phi^{\mathrm{HO}} \end{pmatrix}, \tag{12}$$

can be decomposed as

$$\begin{cases} \phi^{\mathrm{TT}} = \phi^{\mathrm{tur,TT}} + \phi^{\mathrm{vib,TT}} \\ \phi^{\mathrm{HO}} = \phi^{\mathrm{tur,HO}} \end{cases}. \tag{13}$$

We thus have:

$$\phi^{\mathrm{tur}} = \begin{pmatrix} \phi^{\mathrm{tur,TT}} \\ \phi^{\mathrm{tur,HO}} \end{pmatrix}, \quad \phi^{\mathrm{vib}} = \begin{pmatrix} \phi^{\mathrm{vib,TT}} \\ 0 \end{pmatrix}. \tag{14}$$

In the case of CANARY we have checked that there were no noticeable vibration peaks on temporal PSDs for modes above tip and tilt, hence the restriction of the specific treatment to these two modes. Note however that the above equations can be easily extended to the case where vibrations may affect higher orders than tip and tilt, as expected on large structures such as ELTs. In this case, the superscript TT has just to be replaced with LO (for low orders) and the corresponding vectors will contain as many components as needed.

A widely used class of models to describe the phase temporal correlation is the class of vector-valued autoregressive (AR) processes. First studies have shown that an AR model of order 2 (AR2) for each turbulent phase mode leads to good performance, see [32] and in the case of CANARY see [33]. We can thus model the turbulence process as a discrete-time stochastic process in the form

$$\phi_{k+1}^{\mathrm{tur}} = A_1^{\mathrm{tur}} \phi_k^{\mathrm{tur}} + A_2^{\mathrm{tur}} \phi_{k-1}^{\mathrm{tur}} + v_k^{\mathrm{tur}} \tag{15}$$

where A_1^{tur} and A_2^{tur} are diagonal matrices and $v^{\mathrm{tur}} \triangleq \{v_k^{\mathrm{tur}}\}_{k \geq 0}$ is a generating Gaussian white noise.

3.3. Tip-tilt modeling

The vibrations affecting the system can also be modeled by an AR2 [16] (resonant order 2 model). Tip and tilt are treated separately and considered as decoupled. For either tip or tilt components each particular vibration peak will be associated with a scalar vibration model in the form

$$\phi_{k+1}^{\mathrm{vib},\bullet,i} = a_1^{\mathrm{vib},\bullet,i} \phi_k^{\mathrm{vib},\bullet,i} + a_2^{\mathrm{vib},\bullet,i} \phi_{k-1}^{\mathrm{vib},\bullet,i} + v_k^{\mathrm{vib},\bullet,i} \tag{16}$$

where $\phi_k^{\mathrm{vib},\bullet,i}$ is a scalar variable and $\bullet$ stands for tip or tilt. The scalar parameters $a_1^{\mathrm{vib},\bullet,i}$ and $a_2^{\mathrm{vib},\bullet,i}$ depend on damping coefficient and resonant frequency, and $v^{\mathrm{vib},\bullet,i} \triangleq \{v_k^{\mathrm{vib},\bullet,i}\}_{k \geq 0}$ is a white Gaussian excitation noise with variance $\sigma_{v_{\bullet,i}}^2$. The global resulting phase in the pupil can be written as:

$$\phi_k^{\mathrm{TT}} = \phi_k^{\mathrm{tur,TT}} + \phi_k^{\mathrm{vib,TT}} = \phi_k^{\mathrm{tur,TT}} + \begin{pmatrix} \sum_{i=1}^{n_v} \phi_k^{\mathrm{vib,tip},i} \\ \sum_{i=1}^{n_v} \phi_k^{\mathrm{vib,tilt},i} \end{pmatrix} \tag{17}$$

with n_v the total number of vibrations on tip or tilt. As explained in the introduction, these vibration models will be identified on-sky from WFS data (see Section 4.1).

As already mentioned, the extension to vibrations on other modes than tip and tilt is straightforwardly obtained by considering that $\bullet$ stands for all the concerned modes, and with superscripts tip and tilt replaced by a mode number. The right hand term in Eq. (17) will therefore have as many lines as the number of modes affected by vibrations.

3.4. SCAO state-space representation

The state vector can be simply chosen as the concatenation of turbulence and vibration phases at two consecutive time steps:

$$x_k \triangleq \begin{pmatrix} x_k^{\text{tur}} \\ x_k^{\text{vib}} \end{pmatrix}, \quad x_k^{\text{tur}} \triangleq \begin{pmatrix} \phi_k^{\text{tur}} \\ \phi_{k-1}^{\text{tur}} \end{pmatrix}, \quad x_k^{\text{vib}} \triangleq \begin{pmatrix} x_k^{\text{vib,tip}} \\ x_k^{\text{vib,tilt}} \end{pmatrix}, \quad x_k^{\text{vib},\bullet} \triangleq \begin{pmatrix} x_k^{\text{vib},\bullet,1} \\ x_k^{\text{vib},\bullet,2} \\ \vdots \\ x_k^{\text{vib},\bullet,n_v} \end{pmatrix}, \quad x_k^{\text{vib},\bullet,i} \triangleq \begin{pmatrix} \phi_k^{\text{vib},\bullet,i} \\ \phi_{k-1}^{\text{vib},\bullet,i} \end{pmatrix}.$$

(18)

The models presented above can then be re-arranged into a compact state-space model in the form:

$$\begin{cases} x_{k+1} = Ax_k + \Gamma v_k \\ y_k = Cx_k - DNu_{k-2} + w_k \\ \phi_k = C_\phi x_k \end{cases}.$$

(19)

The last line in Eq. (19) expresses the perturbation as a sum of contributors stored in the state x. Both the WFS measurement y_k and the disturbance ϕ_k are thus obtained as outputs of the state model. The other vectors and matrices in the state space representation are defined as

$$v_k = \begin{pmatrix} v_k^{\text{tur}} \\ v_k^{\text{vib}} \end{pmatrix}, \quad A = \begin{pmatrix} A^{\text{tur}} & 0 \\ 0 & A^{\text{vib}} \end{pmatrix}, \quad A^{\text{tur}} = \begin{pmatrix} A_1 & A_2 \\ I & 0 \end{pmatrix}$$

(20)

$$A^{\text{vib}} = \begin{pmatrix} A^{\text{vib,tip}} & 0 \\ 0 & A^{\text{vib,tilt}} \end{pmatrix}, \quad A^{\text{vib},\bullet,i} = \begin{pmatrix} a_1^{\text{vib},\bullet,i} & a_2^{\text{vib},\bullet,i} \\ 1 & 0 \end{pmatrix}$$

$$A^{\text{vib},\bullet} = \begin{pmatrix} A^{\text{vib},\bullet,1} & & 0 \\ & \ddots & \\ 0 & & A^{\text{vib},\bullet,n_v} \end{pmatrix}$$

where $v \triangleq \{v_k\}_{k \geq 0}$ and $w \triangleq \{w_k\}_{k \geq 0}$ are independent Gaussian white noises with variance-covariance matrices $\text{Var}(v_k) \triangleq \Sigma_v$ and $\text{Var}(w_k) \triangleq \Sigma_w$, and where

$$\Gamma \triangleq \begin{pmatrix} \Gamma^{\text{tur}} & 0 \\ 0 & \Gamma^{\text{vib}} \end{pmatrix}, \quad \Gamma^{\text{tur}} = \begin{pmatrix} I \\ 0 \end{pmatrix}, \quad \Gamma^{\text{vib}} = \begin{pmatrix} \Gamma_0 & & 0 \\ & \ddots & \\ 0 & & \Gamma_0 \end{pmatrix}, \quad \Gamma_0 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

(21)

To obtain the perturbation as the sum of turbulence and vibrations, the matrix C_ϕ is defined as

$$C_\phi \triangleq \begin{pmatrix} C_0^{\text{tur}} & C_0^{\text{vib}} \end{pmatrix}$$

(22)

with

$$C_0^{\text{tur}} = \begin{pmatrix} I & 0 \end{pmatrix}, \quad C_0^{\text{vib}} = \begin{pmatrix} c_0 & \dots & c_0 & 0 & \dots & 0 \\ 0 & \dots & 0 & c_0 & \dots & c_0 \\ 0 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & 0 & \dots & 0 \end{pmatrix}, \quad c_0 = \begin{pmatrix} 1 & 0 \end{pmatrix}.$$

(23)

Likewise, to obtain $Cx_k = D\phi_{k-1}$, the matrix C should be defined as

$$C \triangleq D \begin{pmatrix} C_1^{\text{tur}} & C_1^{\text{vib}} \end{pmatrix}.$$

(24)

with

$$C_1^{\text{tur}} = \begin{pmatrix} 0 & I \end{pmatrix}, \quad C_1^{\text{vib}} = \begin{pmatrix} c_1 & \dots & c_1 & 0 & \dots & 0 \\ 0 & \dots & 0 & c_1 & \dots & c_1 \\ 0 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & 0 & \dots & 0 \end{pmatrix}, \quad c_1 = \begin{pmatrix} 0 & 1 \end{pmatrix}. \tag{25}$$

3.5. Kalman filter and control computation

The optimal prediction $\widehat{x}_{k+1|k}$, which enters the optimal control expression in Eq. (5), is computed recursively as an output of the Kalman filter based on the state-space model (19). Since the control criterion is defined on an infinite horizon, we restrict ourselves, without any performance degradation, to the steady-state time invariant version of the Kalman filter [34]. Real-time implementation boils down to the so-called update and prediction equations:

$$\begin{cases} \widehat{x}_{k|k} = \widehat{x}_{k|k-1} + H_\infty(y_k - \widehat{y}_{k|k-1}) = \widehat{x}_{k|k-1} + H_\infty(y_k - C\widehat{x}_{k|k-1} + DNu_{k-2}) \\ \widehat{x}_{k+1|k} = A\widehat{x}_{k|k} \end{cases} \tag{26}$$

where H_∞ is the asymptotic Kalman gain computed off-line as

$$H_\infty \triangleq \Sigma_\infty C^{\text{T}}(C\Sigma_\infty C^{\text{T}} + \Sigma_w)^{-1}, \tag{27}$$

Σ_∞ being the unique solution of the algebraic Riccati equation

$$\Sigma_\infty = A\Sigma_\infty A^{\text{T}} + \Gamma\Sigma_\infty \Gamma^{\text{T}} - A\Sigma_\infty C^{\text{T}}(C\Sigma_\infty C^{\text{T}} + \Sigma_w)^{-1}C\Sigma_\infty A^{\text{T}}. \tag{28}$$

Finally we are able to compute the control vector applied to the DM:

$$u_k = P_u C_\phi \widehat{x}_{k+1|k} \tag{29}$$
$$= P_u(\widehat{\phi}_{k+1|k}^{\text{tur}} + \widehat{\phi}_{k+1|k}^{\text{vib}})$$

where P_u depends on the control criterion: Eq. (6) for residual phase variance and Eq. (10) for residual slopes variance. The control gain that relates the state estimate $\hat{x}_{k|k}$ to the optimal control is thus simply $P_u C_\phi A$.

4. Practical implementation

4.1. Model identification for turbulence and vibrations

We describe in this Section the identification procedures used to derive the parameters of our AR2 models. As explained before, we apply a separate treatment for the TT and the HO:

- TT components are identified thanks to the method presented in [17], which allows identification of a TT turbulent model together with vibration components. This method uses open-loop slopes to compute the mean slopes PSD. The superimposition of several AR2 models is fitted to this PSD: a low frequency AR2 model for turbulence, and n_v vibration components. For each identified vibration the method estimates the resonant frequency, damping coefficient and power (*i.e.* variance of the excitation noise). The number of vibration components n_v to be identified is set to 10 both for tip and tilt. Note that pseudo-open loop slopes could be used as well if residual slopes only were available.

- As mentioned above, HO modes are not affected by vibrations, so that their model can be based on the heuristic formula in [8, 35] derived from the Zernike turbulent temporal

PSD structure [36]. It relates the AR coefficients of the i^{th} mode to wind speed norm and radial order and is an extension of the method proposed in [8]:

$$\begin{cases} \omega_0 = 0.3(n(i)+1)(\overline{V}_{\text{wind}}/D_{\text{pup}})2\pi \\ \omega_r = \omega_0\sqrt{1-\xi^2} \end{cases} \tag{30}$$

$$\begin{cases} a_1 = 2e^{-\xi\omega_0 T}\cos(\omega_r T) \\ a_2 = -e^{-2\xi\omega_0 T} \end{cases} \tag{31}$$

where $\overline{V}_{\text{wind}}$ is the average wind speed norm value in the pupil, D_{pup} is the pupil diameter, $n(i)$ the Zernike radial order of mode i, ξ is the damping factor set to 0.9, T is the sampling period. Note that the model noise variance $\sigma_{v_i}^2$ is derived from the HO covariance matrix deduced from the Fried parameter. Wind speed and Fried parameter are estimated thanks to the LEARN algorithm [31, 37]. The maximum radial order has been chosen equal to $n_{\text{rad}} = 14$.

4.2. Complexity, real-time control and delays

From Eq. (18), the state vector x includes the values of all phase variables at two consecutive time steps. The phase value at each time step is expressed using $n_{\text{rad}}(n_{\text{rad}}+3)/2 = 119$ Zernike modes. When filtering vibrations with 10 vibrations peaks both on tip and tilt, the size of the state x is therefore $n_x = 2 \times (119 + 2 \times 10) = 278$. The number of measurements is equal to $n_y = 36 \times 2 = 72$. We have 52 actuators on the DM and 2 actuators for the TT mirror, so that the size of u is $n_a = 54$ components. To speed computation, the LQG implementation of the DARC RTC has been optimised to remove unnecessary operations (such as products of zero or identity matrices). Consequently, the number of floating multiplications performed in real time is equal to $n_x/2 + n_x^2/2 + 3n_x n_a/2 + n_x n_y \approx 52$k per frame for this configuration.

The Kalman gain computation, which centrally involves solving the Riccati equation (28), is not part of the RTC data pipeline and is performed off-line on a separate interface computer, in our case a standard desktop device. For a state vector with 278 components, solving the Riccati equation took only about 1 to 2 seconds using the efficient doubling algorithm [28].

Both on- and off-line computation times are not a limiting factor for routine operations, and the very flexible general formulation given in Eqs. (26)–(29) makes it advantageous. Of course when moving to more complex systems for VLTs, or even ELTs, both off-line and on-line calculations may become an issue. However various approaches have been recently proposed to overcome this problem by exploiting the particular structure of the operations (linear shift invariance, sparsity, structure...) [38–41], and RTCs' computing power keeps increasing.

Another important practical and theoretical issue is the influence of loop delay on closed-loop performance. In the system model presented in Section 3, we make the standard assumption that the total loop delay between DM control u and WFS output y in the measurement equation (8) is equal to two frames. However, during on-sky operations the actual system loop delay d was measured at around $d = 1.6$ frames for a loop frequency of 150 Hz (the value usually selected in order to achieve adequate signal-to-noise ratio for a 10-12 magnitude star at visible wavelengths). In all the simulations and replays presented in this paper, the system loop delay was therefore set to $d = 1.6$. With integrator control, it is easily seen that lower delay means better disturbance rejection (in addition, as we shall see, it also renders integrator performance considerably less sensitive to changes in the controller gain). The LQG controller is based on an explicit model of the AO system, including delays. In our case, it has been designed using a model with a loop delay of two frames, which simplified its implementation in DARC RTC. This creates a system/model mismatch which may lead to a slight degradation of LQG performance. However, in order to be consistent with on-sky implementation, the value of two frames for the model loop delay in the LQG controller has been kept in all what follows.

5. On-sky results

We present in this Section the first results obtained on-sky using an LQG control on all modes. Taking advantage of the CANARY RTC's flexibility, LQG control (with and without vibration mitigation) and standard integral control were alternatively put on line over successive periods, during which sequences of residual slopes, open-loop slopes and PSF images were recorded. This enabled to compare the performance of these three controllers (LQG with and without vibration compensation and integrator) in the most possible similar on-sky conditions. However, sudden and important variations of the Fried parameter r_0 regularly occur, which makes fair comparisons of performance series often difficult. We have thus preferred to compare performance as a function of r_0, estimated as further explained. These SCAO results have been obtained during the nights of August 28[th] and 29[th] 2012 (time is indicated hereafter in GMT+1, La Palma local time) on the star BD + 10 3697 of magnitude $m_V = 9.7$.

The WFS slope noise estimated on the slope data is of the order of 80 nm rms, expressed as an optical path difference at the edge of a WFS subaperture. This noise is mainly a consequence of photon noise at the level of the WFS subaperture images detection, since we use a low detector noise EMCCD ANDOR camera. Its impact on performance is difficult to estimate precisely and depends on the choice of the controller. We still estimate in our case that an upper bound of its contribution is 10% of the global residual phase variance in H band.

5.1. Tuning parameters for control

5.1.1. Integrator

The control computation for the integrator is given by:

$$u_k = u_{k-1} + g_{\mathrm{int}} M_{\mathrm{com}} y_k \tag{32}$$

where g_{int} is the global feedback gain (i.e., the same gain was used for all modes). This gain has been optimized on-sky for each night of the CANARY runs. In Fig. 1, we give an example of on-sky determination of the gain. We can see that performance is not very sensitive to the gain value. For this particular night, the gain was set to 0.6. This flatness of the on-sky performance sensitivity seems counterintuitive, since one would normally expect a parabolic-like performance curve. However, such parabolic performance would correspond to a total loop delay of two frames, whereas for $d = 1.6$ the expected performance becomes essentially flat (more on this in Section 5.3).
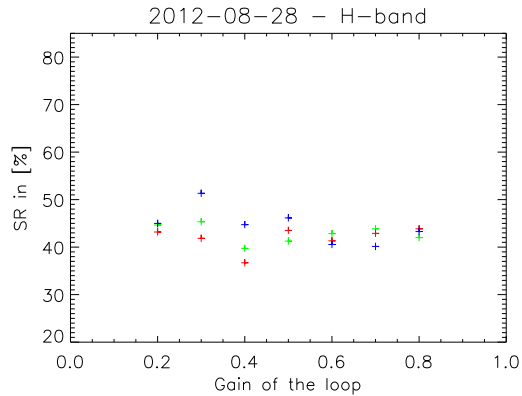
Fig. 1. On-sky optimization of the integrator gain during the observation done the 28th of August 2012. Different colors represent three gain tuning sequences at different times of the night.

5.1.2. LQG

A convenient tuning parameter in LQG control is the measurement noise variance Σ_w taken in the model. Performance is generally improved when setting a noise level higher than the physical noise values so as to absorb various model uncertainties inherent to experimental operation [9, 20, 33]. Prior to starting on-sky operation, this relationship was confirmed through investigation using the CANARY telescope simulator in similar noise conditions to that expected on sky. An amplification of the noise variance by a factor of 2 improved performance and thus this tuning value was used for subsequent on-sky tests.

5.2. *Performance results*

We first present overall performance in terms of Infra-Red (IR) Strehl Ratio (SR), and then focus more specifically on the TT correction.

5.2.1. Overall performance

The SRs are estimated on long exposure images (typically 3 sec) recorded with CAMICAz in H or K-band. SR are computed with a particular care for what concerns dead pixel correction, background subtraction and fine re-centering. For each image we associate a Fried parameter value r_0 estimated from the 3 NGS off-axis WFS open-loop slope data (typically datasets of 2048 samples at 150 Hz, hence $\approx$ 14 seconds).
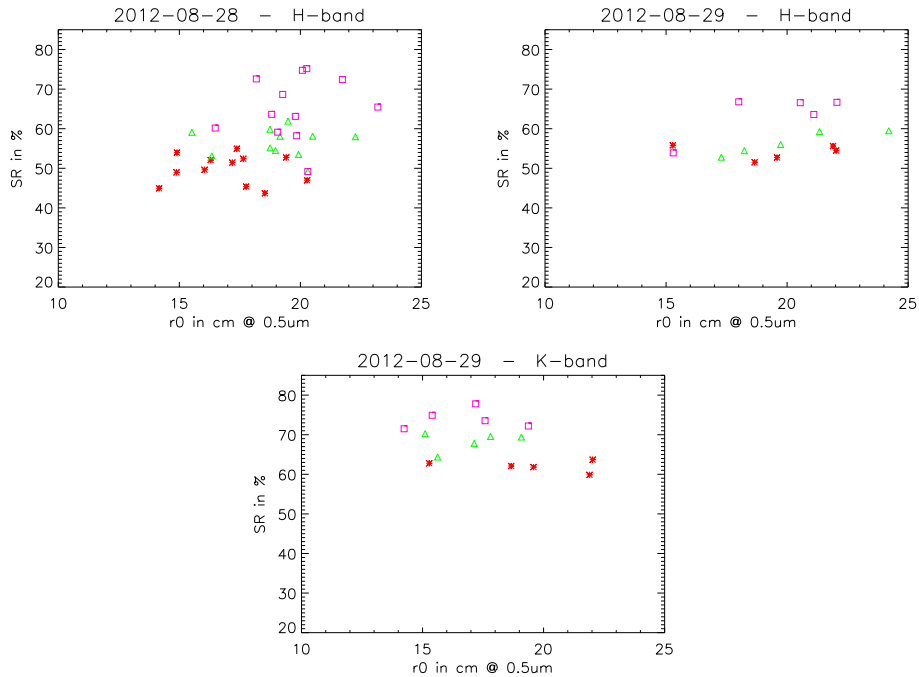
Fig. 2. Strehl ratios as a function of the Fried parameter for integrator and LQG controllers. Top left: SR in H-band recorded on Aug. 28th 2012. Top right: SR in H-band recorded on Aug. 29th 2012. Bottom: SR in K-band recorded on Aug. 29th 2012. Colors: red stars = integrator control, green triangle = LQG control without vibration mitigation, and magenta square = LQG control with vibration mitigation.

Figure 2 presents the SRs obtained on August 28th 2012 in H-band and on August 29th 2012 in H and K-band. We compare in each case integrator control (gain set to 0.6), LQG control with no specific handling of vibrations and LQG control with vibration filtering. Despite strong dispersion in seeing conditions, we clearly see that LQG with vibration mitigation indeed brings a gain in performance compared to LQG without vibration mitigation. Moreover, the LQG controller behaves as expected, globally better than integrator, in terms of performance and stability.

The turbulent priors used in the model building for LQG control ($\overline{V}$, r_0, Σ_w) are identified thanks to the LEARN algorithm [31, 37] which uses an open-loop WFS data sequence of 5000 samples. At a sampling frequency of 150 Hz, this corresponds to 33 seconds of acquisition. These data sets are also used for the identification of the TT models. We record a new set of such open-loop data every time we change configuration or asterism on the sky to re-estimate all priors and parameters. LQG matrices were updated in this fashion about every hour.

The fact that performance is not very sensitive to r_0 may seem surprising at first glance. However a gross estimate of error contributors for the integrator shows that more than half the residual phase variance is related to non-seeing dependent terms (Fig. 3). These terms are mainly due to uncorrected bench static errors and to vibrations.
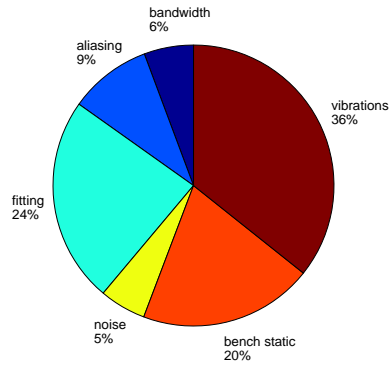
Fig. 3. Example of integrator control relative error budget in H band for a data sample on August 29th 2012 with $r_0 = 21$ cm.

Figures 4 and 5 show some examples of Point Spread Functions (PSFs) obtained with the 3 control strategies (integrator, LQG with and without vibration mitigation). As the controller becomes more efficient, SR improves, the PSF becomes more symmetric and first Airy ring more contrasted.
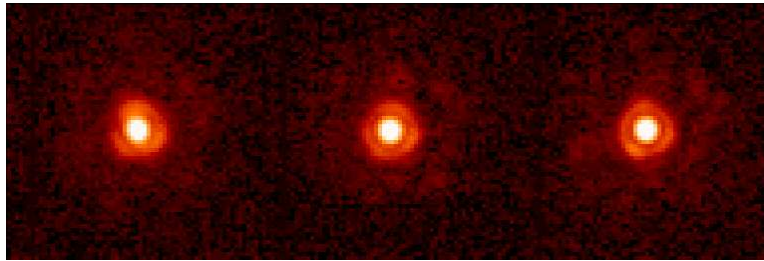

Fig. 4. PSFs recorded in K-band on August 29[th] 2012 observing BD + 10 3697. Left: integrator (SR = 62.1% @ $r_0 = 16.1$ cm). Center: LQG (SR = 70.3% @ $r_0 = 15.1$ cm). Right: LQG with vibration mitigation (SR = 77.8% @ $r_0 = 17.2$ cm).
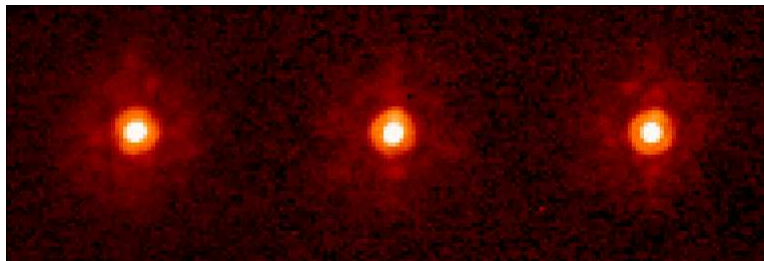
Fig. 5. PSFs recorded in H-band on August 29[th] 2012 observing BD + 10 3697. Left: integrator (SR = 50.9% @ $r_0 = 19.4$ cm). Center: LQG (SR = 61.9% @ $r_0 = 19.5$ cm). Right: LQG with vibration mitigation (SR = 68.7% @ $r_0 = 19.3$ cm).

5.3. Performance analysis in replay mode

These on-sky comparisons between LQG and integrator controller raise two important questions. First, would LQG retain its edge when compared with a properly tuned Optimized Modal Gain Integrator (OMGI), where different feedback gains are used for different turbulent modes in lieu of the global gain g_{int}? A second relevant question is whether the gap in performance between LQG and integrator is mainly due to vibration peaks and their particular localization, in which case the said performance gap might just as well vanish altogether for slightly different on-sky conditions.

In order to address these issues, we exploited the fact that during CANARY experiments a sequence of open-loop WFS slopes is systematically recorded every time a control is tested. To evaluate the performance that any other controller would have achieved on-sky at this particular moment, it therefore suffices to feed this open-loop WFS data sequence in a filter that emulates its closed-loop rejection transfer function. This produces sequences of residual slopes, from which we extracted residual phase trajectories by applying a minimum-variance slopes-to-Zernike modes reconstructor. This so-called 'replay mode' enabled to replicate with excellent accuracy the actual on-sky performance of both LQG and integral controllers. Furthermore, replay mode simulations reveal that integrator performance is highly sensitive to changes in the system loop delay d. Figure 6 presents the performance variation as a function of the gain, obtained by replaying the data in two cases: $d = 1.6$ frame and $d = 2$ frame delay. The same behaviour has been obtained on all data sets: the 1.6 frame delay curve is essentially flat, which is consistent with on-sky results in Fig. 1.
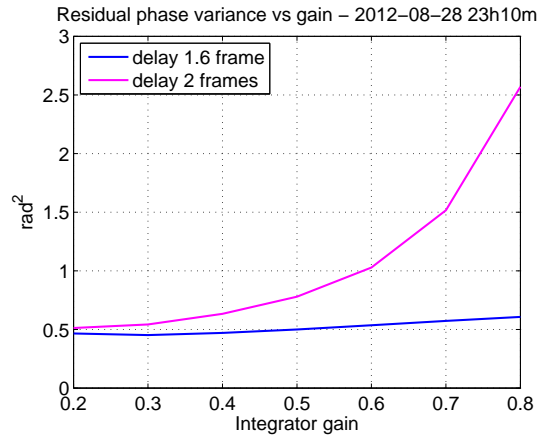

Fig. 6. Residual phase variance in rad^2 for integrator as a function of global feedback gain g_{int} obtained by replaying the open-loop on-sky data of the 28th August at 23:10. Two different loop delays are considered: $d = 1.6$ (blue) and $d = 2$ (magenta).

5.3.1. Optimized modal gain integrator performance in replay

As already mentioned, the same gain has been used for all modes in on-sky integrator tests. To explore in replay the benefit of choosing a different integrator gain for each mode, Zernike modes were reconstructed from open-loop on-sky WFS data; we then applied to each sequence the integrator rejection transfer functions corresponding to values of g_{int} ranging from 0.2 to 0.8 and draw performance curves in terms of residual phase variance (in rad^2) for each reconstructed mode. This detailed modal analysis showed that, due to the presence of vibrations affecting only tip and tilt and to the moderate WFS noise level, best OMGI performance could

be achieved by using only two different gains: a lower value g_{TT} for tip and tilt modes and a higher value g_{HO} set identical for all higher-order modes.

Figure 7 shows representative performance levels (in rad^2) obtained on the same open-loop data-set (August 29th at 22:07) for this gain tuning strategy. The blue curve is the performance of the integral controller when tip and tilt gains are set to $g_{TT} = 0.3$ while the higher orders gain varies from $g_{HO} = 0.2$ to $g_{HO} = 0.8$. The red curve corresponds to another tuning, where the higher orders gain is set to $g_{HO} = 0.6$ while the tip and tilt gain g_{TT} varies from 0.2 to 0.8. Note that for $g_{TT} = 0.3$ the performance curve exhibits a minimum around $g_{HO} = 0.6$ and is almost flat for $0.5 \leq g_{HO} \leq 0.8$. Thus, a finely tuned OMGI controller would give roughly the same performance as $g_{TT} = 0.4$ and $g_{HO} = 0.6$ – a quasi-optimal setting which corresponds to the lowest point on the red curve. The on-sky setting for the global integrator gain during this night was $g_{int} = 0.6$, which corresponds to a quasi-optimal choice of g_{HO} but a markedly sub-optimal choice of g_{TT}. It also appears that using this setting instead of the optimal tuning would increase the residual variance by about 5%, whereas LQG residual phase variance (in magenta) is 25% lower than this best integrator performance. This trend was confirmed on the many data sets recorded during different nights.
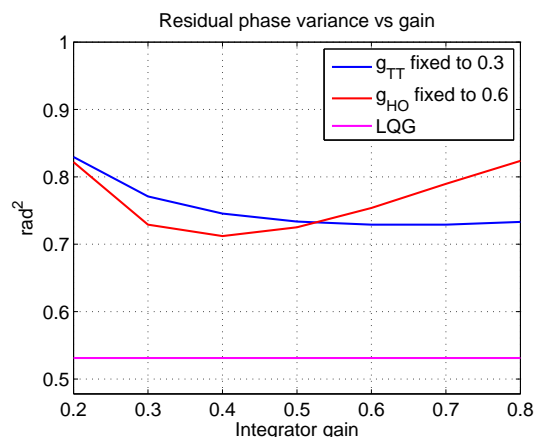

Fig. 7. Integrator performance in rad^2 obtained in replay on the open-loop on-sky data of the 29th August-22:23 for different gain values. Blue: tip and tilt gain $g_{TT} = 0.3$ and higher orders gain $0.2 \leq g_{HO} \leq 0.8$. Red: tip and tilt gain $0.2 \leq g_{TT} \leq 0.8$ and $g_{HO} = 0.6$. Pink: LQG performance.

5.3.2. Vibration mitigation/amplification and performance

Comparing on-sky performance of different controllers is difficult because disturbance conditions vary over very short time periods. Replay mode offers a neat way to circumvent this difficulty, since it enables to compare the performance that different controllers would achieve when confronted to exactly the same disturbance trajectory. To assess the impact of vibration on closed-loop integrator and LQG performance, we thus present here results obtained by replaying an open-loop sequence of WFS measurements acquired on the night of August 29th with the 3 controllers (integrator, LQG with and without vibration compensation) that had been tuned and tested on-sky at that point in time. For LQG control, this means that we have used here the matrices that have been computed during on-sky operations.

From this sequence of on-sky open-loop slopes and these 3 sequences of residual slopes obtained in replay (one for each controller), we extracted tip and tilt modes trajectories using a

modal reconstructor (pseudo-inverse of *D*) and computed the corresponding 'tip plus tilt' cumulated PSDs (i.e. the sum of tip and tilt cumulated PSDs at each frequency). These cumulated PSDs, which thus provide for each controller a frequency-by-frequency total tip and tilt residual variance, are presented in Fig. 8.
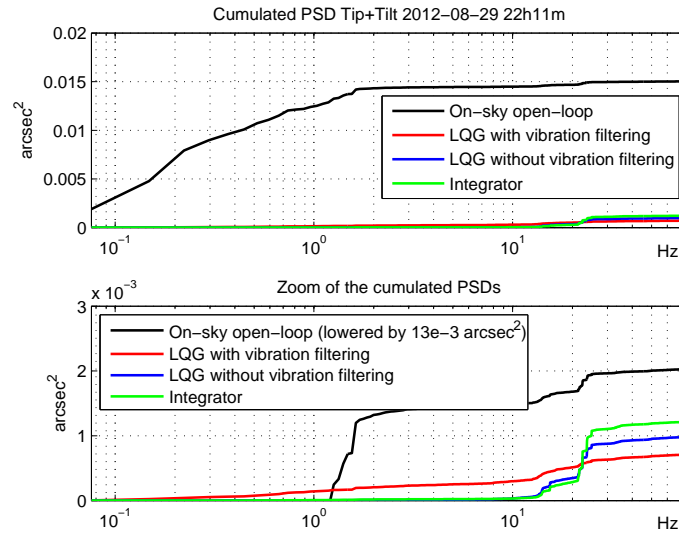

Fig. 8. Cumulated PSDs for open-loop and residual tip plus tilt disturbance and residuals computed from WFS data recorded on-sky on Aug. 29th 2012. Tip and tilt are evaluated in equivalent on-sky angle, expressed in arcsec, hence PSDs in $arcsec^2$. (Open-loop data on bottom plot has been lowered by 0.013 $arcsec^2$ to increase vertical resolution.)

The open loop curve presents several steps indicating vibration peaks at different frequencies. At frequencies up to a few hertz they are strong but have very limited impact on performance since all controllers have a good rejection (see Section 6.2). The last vibration peak, located at 22 Hz, is more critical since it is amplified by the integrator control (see Section 6.2 for details) and not corrected by standard LQG control. However, LQG with dedicated vibration identification and filtering mitigates this contribution. On these data we estimate that the 22 Hz vibration contribution amounts, for the integrator control, to a loss of SR in K band by a factor 0.9 .

For results such as the ones in Fig. 8, it is clear that the vibration peak affecting tip and tilt modes around 22 Hz is amplified by the integrator and mitigated by the LQG controller. To make a quantitative assessment of the impact on overall performance for this particular on-sky dataset, we have evaluated the corresponding cumulated open-loop and residual tip plus tilt DSPs over the frequency range [10 Hz − 75 Hz]. The step for the open-loop PSD (tip+tilt) is $5\,10^{-4}$ $arcsec^2$, to be compared with $3.4\,10^{-4}$ for LQG and $10.9\,10^{-4}$ for integrator, which clearly amplifies the vibrations. On this example (fairly representative of our on-sky results), the tip+tilt residual variance of the integrator is 107% higher than LQG. We could expect an OMGI controller to give better results, and we have thus evaluated its performance in replay. However, even with a finely tuned OMGI controller (in our simulation, $g_{TT} = 0.3$ and $g_{HO} = 0.7$), the residual variance would still remain 43% higher than with LQG.

While many real AO systems exhibits vibration spectra spanning a wide frequency range (see

for example [12]), we may also consider ideal 'vibration-free' situations, i.e. where perturbations exhibiting no vibration peak in any 'critical' frequency range. Would high-performance SCAO controllers still improve performance? We have thus conducted tests to give an idea of what could be gained with an LQG controller on such data for which the vibrations within the integrator overshoot zone have been removed. This was achieved by modifying on-sky data sets to obtain filtered data without vibrations in the range $10 - 40$ Hz. An example of such spectra, corresponding to the data of August 29^{th} 2012 at 22:02, is given in Fig. 9.
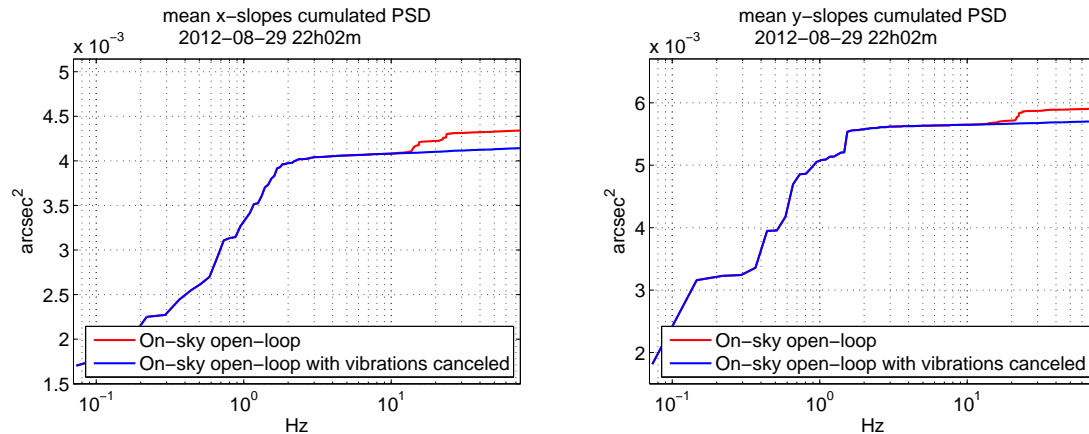
Fig. 9. On-sky original x- and y-slopes CANARY data (red) and a filtered version (blue) where vibrations have been removed in the range 10-40 Hz. Data set of August 29^{th} 2012 at 22:02.

We have evaluated the performance obtained by replaying the loop on these data for the LQG controller – with TT model based on Eqs. (30)–(31) – and for the integrator controller with a fixed gain of 0.7 (the best tuning in replay). The LQG residual variance is around 6% better, thus keeping slightly ahead. Depending on the processed data set, this percentage varies between 5 and 10 %. Note that some improvement could be gained by designing the LQG controller with a system loop delay of 1.6.

6. LQG vs. integrator: frequency-domain and robustness analysis

This Section presents a comprehensive comparative analysis of LQG and integrator performance from a frequency-domain perspective. We seek to provide answers to the following questions: 1/ Do the disturbance models used for LQG control accurately match the turbulence and vibration modes' PSDs reconstructed from on-sky data? 2/ Is the quality of LQG correction for these modes, as measured by their modal rejection transfer functions, significantly better than with an integrator? 3/ Is LQG control robust to modeling uncertainties?

These three issues are in fact closely intertwined, because of the fundamental performance/stability trade-off inherent to any control loop: on the one hand, efficient disturbance rejection requires high feedback gains; on the other, (too) high feedback gains render the loop's behavior unstable or at least too sensitive to modeling uncertainties. In order to negotiate a sensible trade-off between these inherently conflicting requirements, the controller gain thus needs to be finely adjusted according to the disturbance's PSD and also to the measurement signal-to-noise ratio.

6.1. Model relevance for turbulence and vibration

We discuss here the relevance of the models of high-order (Section 6.1.1) and tip-tilt modes (Section 6.1.2) used for LQG control by comparing their temporal PSDs to those of on-sky data.

6.1.1. Model matching for high orders

Figure 10 compares the temporal PSDs of the AR2 models used as priors for turbulence with the experimental PSDs computed on reconstructed Zernike coefficients (Z4 and Z28).

We recall that the high order AR2 models are based on the physical description in Eqs. (30)–(31), with a cut-off frequency proportional to radial order and wind norm, and an energy related to Kolmogorov spectrum (see Section 4.1). The prior temporal PSDs of these models (with wind norm and r_0 estimated on the data) are plotted in Fig. 10. The experimental PSDs have large dynamics, a high level at low frequency (of course larger for low orders), a sharp decrease after the cut-off frequency and a noise floor at high frequency. We see that the cut-off frequency increases as expected with the radial order of the Zernike polynomial [36]. The good agreement between modeled and measured PSDs (even if the cut-off frequencies on left column could be better adjusted) validates the model structure and the parameter estimation. On these curves, the turbulent AR2 models of course do not include the high-frequency noise plateau due to the measurement noise.
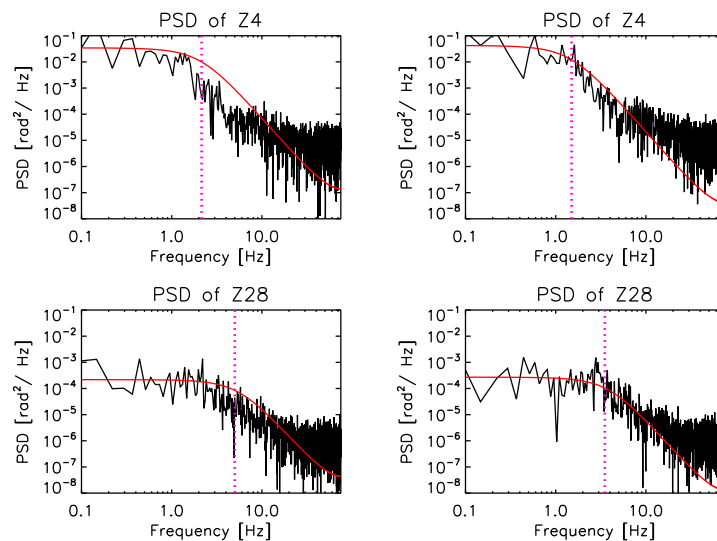

Fig. 10. PSDs of high orders Zernike Z4 and Z28. Black = Regularized reconstruction from open-loop data, red = AR2 model. The theoretical cut-off frequency is indicated by a magenta vertical dotted line. The data have been recorded the August 29[th] 2012. Left: slopes data at 22h07m32s ($\overline{V} = 10.3$m/s). Right: slopes data at 23h07m07s ($\overline{V} = 7.8$m/s). Top: Zernike number 4. Bottom: Zernike number 28. The AR2 models have been identified on the sequence of open-loop data at: left: 21h53m26s and right: 22h54m21s.

6.1.2. Model matching for tip and tilt

We now compare the PSDs of on-sky x and y mean slopes with the model identified from this data set as a sum of AR2 processes. Recall that in the turbulent model, the vibration components and the noise plateau due to measurement noise are identified in the frequency domain

with the method presented in [17]. Figure 11 shows two examples of fits provided by the identification procedure. A good agreement is found both for the smooth turbulence PSDs and for the vibration peaks thanks to the flexibility provided by the adaptation of a damping coefficient for each component. On Fig. 10 (right), a different tuning of the parameter identification algorithm could have led to an even better fit around 15 Hz. However, during all on-sky tests, we retained the same tuning of the identification algorithm.
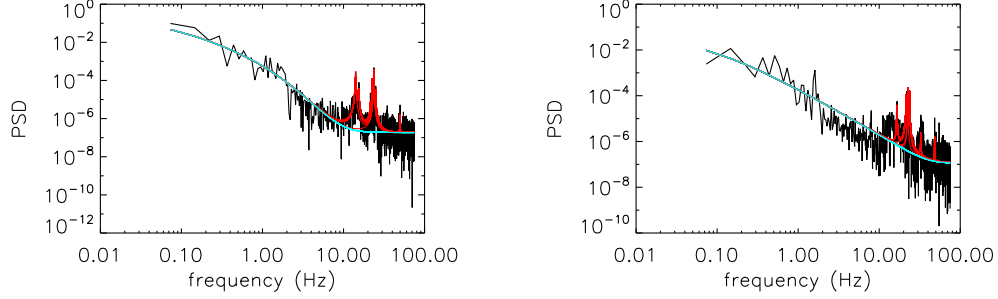
Fig. 11. Identification of vibrations on Tip (left) and Tilt (right) and fit of the AR2 model for vibration filtering in the perturbation modeling.

6.2. Rejection transfer functions

In the sequel, we denote by $E^{\text{tur}}(z)$ the AO loop's disturbance rejection transfer function (a.k.a. the discrete-time feedback loop's sensitivity function), which is the transfer function between the disturbance ϕ and the residual phase ϕ^{res}. To compute $E^{\text{tur}}(z)$, one needs to take into account the value of the system loop delay, that is $d = 1.6$ frame. In practice, a total loop delay $1 < d \leq 2$ means that the control u_k is not applied between WFS sampling instants $t = (k+1)T$ and $t = (k+2)kT$, but instead between $t = (k+d-1)T$ and $t = (k+d)T$. Thus, the correction phase ϕ_k^{cor} seen by the WFS during the sampling interval $[(k-1)T, kT[$ is not equal to Nu_{k-1} but to

$$\phi_k^{\text{cor}} = N\left((d-1)u_k + (2-d)u_{k-1}\right), \tag{33}$$

so that the measurement equation that needs to be accounted for in the computation of $E^{\text{tur}}(z)$ is

$$y_k = D\phi_{k-1} - DN\left((d-1)u_{k-1} + (2-d)u_{k-2}\right) + w_k. \tag{34}$$

Using this formula, it is immediately checked that the rejection transfer function is given by

$$E^{\text{tur}}(z) = \left(I + \left((d-1)z^{-1} + (2-d)z^{-2}\right)NG(z)D\right)^{-1}, \tag{35}$$

where G is the controller transfer function (i.e. the open-loop transfer function from residual slopes to control inputs). To characterize the effect of the rejection transfer function on a particular Zernike mode, it suffices to consider the corresponding diagonal term of the sensitivity function.

From Eq. (32), the integrator controller transfer function is

$$G(z) = \left(1 - z^{-1}\right)^{-1} g_{\text{int}} M_{\text{com}}. \tag{36}$$

Because $NM_{\text{com}}D \approx I$, the matrix $E^{\text{tur}}(z)$ is approximately diagonal, with the same rejection transfer function for all Zernike modes.

For LQG control, from Eqs. (26) and (29), we get:

$$G(z) = \left(I - z^{-1}A(I - H_\infty C) + z^{-2}AH_\infty DNP_uC_\phi\right)^{-1} AH_\infty. \tag{37}$$

6.2.1. Rejection transfer functions for high orders modes

In this Section, we compare the rejection transfer functions of integrator and of LQG for the two Zernike modes Z4 and Z28. Because of a fundamental result in control theory known as Bode's integral theorem, which in its discrete-time version applies to all feedback loops with delay $d \geq 1$ (see e.g. [42], page 79), increasing rejection for some particular frequency will be repaid in kind by disturbance amplification elsewhere (the so-called 'water-bed effect'). As a consequence, it is easily checked that for feedback loops with reasonably good signal-to-noise ratio, the gain of an optimal minimum-variance controller should at low-and-medium frequency be roughly inversely proportional to the square root of the PSD of the perturbation while remaining bounded at high frequencies so as not to amplify measurement noise. This rule of thumb provides a convenient way to gauge, by simple visual inspection, whether or not a particular controller provides optimal or near-optimal rejection [11].

Figure 12 shows the AR2 model temporal PSDs and the rejection transfer functions. As expected, the PSD cut-off frequency increases according to the radial order [36]. Note also that LQG's rejection exhibits a low-frequency 'plateau' which closely matches the turbulence's PSD, resulting in more efficient correction.
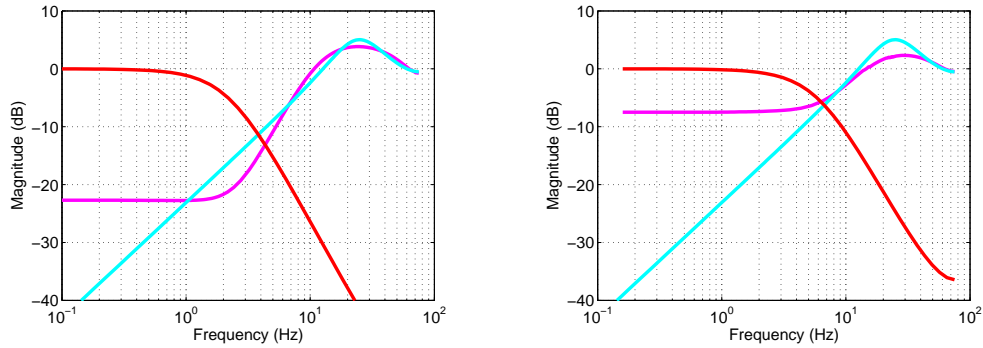

Fig. 12. Left: Rejection transfer functions on Z4. Right: Rejection transfer functions on Z28 (observing night 29 August 2012). Blue for integrator control and Magenta for LQG. The red curve corresponds to the normalized square root of the PSD of the perturbation's model.

6.2.2. Rejection transfer functions for tip and tilt

Figures 13 and 14 show the rejection transfer functions of the TT modes for LQG control and integrator together with the theoretical PSD of the global AR2 TT model (in red), which is the sum of 11 AR2 processes, one for the TT turbulent component and 10 AR2 for identified vibration peaks. As explained above, the integrator has a high rejection at low frequency, but the rejection degrades rapidly and presents even an overshoot centered at about 20 Hz. On CANARY, as we have shown in Section 6.1.2, there are 2 energetic vibration peaks near this overshoot frequency which significantly degrades the performance. LQG controller with vibration filtering reject those vibrations peaks, leading to a much more efficient correction.
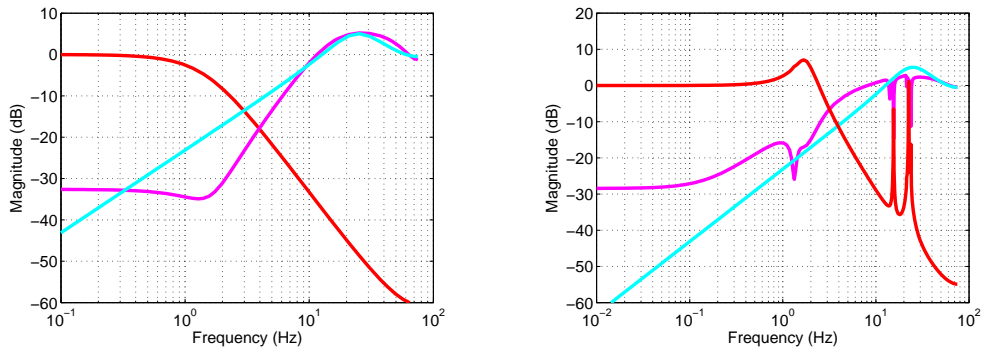
Fig. 13. Rejection transfer functions on tip. Left: without vibration filtering. Right: with vibration filtering. (observing night 29 August 2012). Blue: integrator control. Magenta: LQG. The red curve corresponds to the normalised square root of the PSD of the model perturbation.
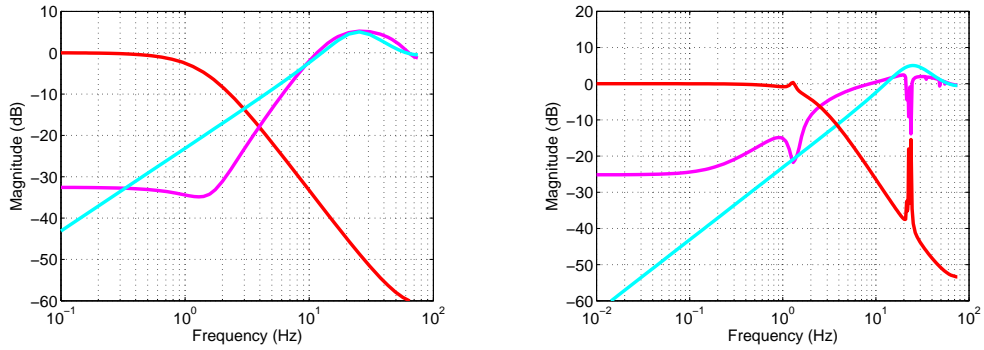

Fig. 14. Rejection transfer functions on tilt. Left: without vibration filtering. Right: with vibration filtering. (observing night 29 August 2012). Blue: integrator control. Magenta: LQG. The red curve corresponds to the normalised square root of the PSD of the model perturbation.

6.3. *Robustness stability margins*

An important issue for all feedback systems is closed-loop stability. For the LQG AO system, closed-loop stability means that the matrix-valued rejection transfer function $E^{tur}(z)$ defined by Eq. (35) should be stable, i.e. should have all its poles (all the values of z for which $E^{tur}(z)$ is not defined) inside the unit circle of the complex plane. Because LQG is a minimum-variance controller, closed-loop stability is guaranteed at least for the nominal value of $E^{tur}(z)$. In addition, an important insight into robust stability – does the closed-loop system remain stable when parameters deviate from their nominal value? – can be derived from the frequency responses of the modal rejection transfer functions (including their phases).

This is achieved by noting that each diagonal term $E_i^{tur}(z)$ of $E^{tur}(z)$ can be viewed [7] as the disturbance rejection transfer function of an equivalent monovariable feedback loop with loop transfer function $L_i(z)$, unitary feedback and additive output disturbance. Thus, $L_i(z)$ can

be retrieved from $E_i^{\text{tur}}(z)$ using the identity:

$$E_i^{\text{tur}}(z) = \frac{1}{1 + L_i(z)}. \tag{38}$$

Under the assumption that $L_i(z)$ is a so-called 'regular' loop transfer function (i.e., without zeros or poles outside the unit circle), the Nyquist stability criterion requires that in order for $E_i^{\text{tur}}(z)$ to be stable, the Nyquist plot of $L_i(z)$ must remain at the right of the critical point -1 for all positive frequencies. The next step is to evaluate the associated gain, phase and delay margins. Thus, denoting as $\theta_{p,1}, \ldots, \theta_{p,m_i}$ the normalized frequencies at which $\arg(L_i(e^{j\theta})) = \pi$, the largest extra gain which can be added without destabilizing the loop is given by the gain margin $M_{G,i}$ defined through

$$M_{G,i} = \min_{k=1,\ldots,m_i} \frac{1}{|L_i(e^{j\theta_{p,k}})|}. \tag{39}$$

Likewise, denoting as $\theta_{c,1}, \ldots, \theta_{c,n_i}$ the normalized frequencies at which $|L_i(e^{j\theta})| = 1$, the largest negative phase shift and the largest extra delay which can be added without destabilizing the loop are given respectively by the phase margin $M_{\phi,i}$ and the delay margin $M_{D,i}$ defined through

$$M_{\phi,i} = \min_{k=1,\ldots,n_i} \left(\arg(L_i(e^{j\theta_{c,k}})) - \pi \right) \tag{40}$$

$$M_{D,i} = \min_{k=1,\ldots,n_i} \frac{\arg(L_i(e^{j\theta_{c,k}})) - \pi}{\theta_{c,k}} \tag{41}$$
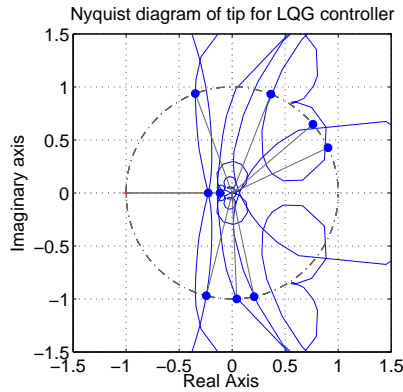

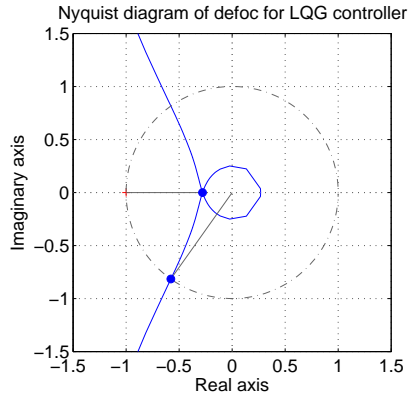Fig. 15. Nyquist plot of LQG loop transfer function, tip

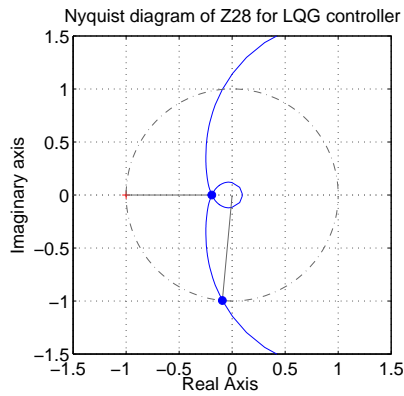Fig. 16. Nyquist plot of LQG loop transfer function, defocus

Fig. 17. Nyquist plot of LQG loop transfer function, Zernike mode 28

Figures 15–17 show the Nyquist plots of $L_i(z)$ for 3 representative Zernike modes (tip, defocus Z3 and Z28) for a system loop delay of 1.6 frames. These Nyquist plots keep clear of the critical point -1, leading to good robustness margins. Table 1 presents the gain, phase and delay margins for 7 different Zernike modes (low, medium and high-order). In order to provide a point of comparison, the table also gives the gain, phase and delay margins for the integrator controller with the standard default setting $g = 0.5$. This suggests excellent robustness of LQG, especially since in this particular case the Kalman filter and the control gain have been computed for an incorrect total loop delay of two frames.

Table 1. LQG and integrator rejection rejection transfer function's gain, phase and delay margins for different Zernike modes

Controller & mode	gain margin	phase margin	delay margin (frames)
LQG Z2 (tip)	12.9 dB	75.9°	1.61
LQG Z3 (tilt)	12 dB	69.4°	1.77
LQG Z4 (defocus)	11.1 dB	54.7°	1.79
LQG Z5 (astigmatism)	12.4 dB	58.2°	2.21
LQG Z7 (coma)	11.5 dB	58.7°	2
LQG Z14	13.3 dB	67.8°	2.76
LQG Z28	14.3 dB	84.7°	3.91
integrator, $g = 0.5$ (all modes)	10.5 dB	59°	2.1

7. Conclusion

We have presented in this paper the first SCAO results obtained on-sky with LQG control on all modes. These results, that include vibration mitigation on tip and tilt, have been obtained on the CANARY demonstrator installed at the William Herschel Telescope on the Roques de Los Muchachos in La Palma (Canaries Islands, Spain). A detailed performance analysis is performed and comparison is conducted with the integral action controller. Note that preliminary LQG results obtained in MOAO with CANARY have been presented in [23].

LQG control gives better on-sky performance than a single gain integrator controller, and proved to be reliable with perfect loop stability and robust performance. Like most AO systems, the CANARY demonstrator is submitted to several vibrations affecting the tip and tilt modes. Using the frequency domain identification method presented in [17], LQG with vibration filtering leads to significant additional improvement in performance. This validates the tip-tilt control strategy adopted on the European extrasolar planet finder SPHERE at the VLT. Extension to vibration mitigation on other modes is straightforward, as indicated in the formulation presentation.

Replays on on-sky data have allowed to test an Optimized Modal Gain Integrator (OMGI). We show that LQG control still leads to better performance. Besides, even in the absence of vibrations in the overshoot frequency range of the integrator, LQG control keeps a slight advantage.

The designed LQG controller relies on a simple and physical state model adapted to the disturbance's PSD which leads to a better compromise between perturbation rejection and noise propagation in the loop. The selected second order autoregressive perturbation models are in good agreement with CANARY on-sky data. Performance is also analysed through evaluation of rejection transfer functions (RTF) and stability margins. Important attenuations in the RTFs confirm the ability to efficiently damp vibratory components. As for stability, excellent gain, phase and delay margins have been evaluated on the on-sky LQG controllers.

This allows to confidently envision more informative models. We could for instance account for wind direction since estimation of wind profiles in norm and direction is achievable with ever better accuracy [43–46]. This may be of particular interest in the context of WFAO [47]. Robustness in stability does not mean robustness in performance, and sensitivity to errors in wind direction estimation thus also needs to be investigated.

Acknowledgments