

Primary mirror control for large segmented telescopes: combining high performance with robustness

Gert Witvoet^a, Niek Doelman^{a,b}, and Remco den Breeje^a

^aTNO Technical Sciences, Optomechatronics department, Delft, The Netherlands

^bLeiden Observatory, Leiden, The Netherlands

ABSTRACT

Future large telescopes, such as E-ELT and TMT, will need feedback control of the thousands of actuators underneath their segmented primary mirrors (M1). Differences in actuator dynamics and spatially and temporally changing disturbances make it extremely difficult to formulate classical controllers which are both sufficiently robust and highly performing. Therefore, TNO has developed and tested a control approach, in which the actual system response is quickly measured, disturbances are continuously estimated and the controller is adapted in real-time. The algorithm is tested on an actual M1-relevant setup, in which it converges to a sub-nm optimum within a few minutes, keeps track of changing disturbances and shows its reliability over multiple days.

Keywords: E-ELT, segmented mirror, position actuator, feedback control, adaptive control

1. INTRODUCTION

The pioneering work on the Keck¹ telescopes paved the way for the use of segmented primary mirrors for optical telescopes. Its example has later been followed by the Hobby-Eberly Telescope and the Southern African Large Telescope (although these telescopes have a fixed declination, whereas Keck has movable azimuth and declination) and the Gran Telescopio Canarias. The enormous sizes of tomorrow's most prominent telescopes, such as the Thirty Meter Telescope (TMT) and the European Extremely Large Telescope (E-ELT), requires them to use a similar segmented primary mirror concept.

In a segmented mirror concept each individual primary mirror (M1) segment must be actively controlled, e.g. in rigid-body piston, tip and tilt motions, in order to create the desired mirror surface shape with sufficient smoothness, even in the presence of disturbances, such as non-stationary structural vibrations and wind loads. To this end each segment is equipped with three dedicated single-DOF actuators, which require extremely high accuracies down to the nanometer level (in case of E-ELT²). However, for larger telescopes this co-phasing challenge becomes increasingly difficult for multiple reasons.


First, dynamics will have a bigger influence, in the sense that each active control loop will 'see' somewhat different servo dynamics.³ On the one hand this is caused by the segment assembly itself, due to mutual variations in e.g. segment geometry and mass and manufacturing tolerances in segment supports. On the other hand for larger and heavier telescopes the telescope structure and its mass and stiffness distribution will play a bigger role; segments and their control systems on e.g. the edge of the primary mirror will experience a different back structure than the ones in the center, again yielding differences in servo dynamics.

Second, for large telescopes the disturbances on each segment vary over the primary mirror. Not only will the non-stationary structural vibrations differ from segment to segment, partly due to spatial differences in the back structure dynamics, but wind loads also vary heavily spatially. Segments high up in the telescope will simply capture more wind forces than the ones down under, hence some control systems will need to achieve much more disturbance attenuation than others. Moreover, the disturbances will also change temporally, both due to changes in telescope azimuth and declination, and due to changing weather conditions.


Third, large segmented primary mirrors induce complexity simply by numbers. In the E-ELT case there will be a total of 2394 piston-tip-tilt actuators, each of them actively controlled. As mentioned, each of these

Further author information:


G. Witvoet: E-mail: Gert.Witvoet@tno.nl, Telephone: +31 (0) 888 666 483



(a) Picture of a PACT prototype.
Photo: TNO / Fred Kamphues.



(b) Schematic of PACT



(c) Drawing of PACT in its test setup

Figure 1: Overview of the TNO PACT hardware and M1-relevant test setup.

2394 control loops will experience somewhat different dynamics, and will have different disturbance attenuation requirements to reach the desired accuracy, which will even change over time. When using classical control strategies, this imposes a dilemma. Designing one robust controller for all actuators will inevitably sacrifice performance, while designing 2394 different controllers for all possible disturbances will be extremely laborious or even infeasible. Moreover, such a classical approach is typically time-invariant and will not be able to track the changing disturbances.

Achieving a smooth optical surface at all locations at all time is thus a major challenge for telescopes with larger segmented M1, since it boils down to the classical trade-off in control theory between robustness and optimal performance on each control loop. To overcome this trade-off for mirror segment actuators, the Netherlands Organisation for Applied Scientific Research (TNO) has developed and tested an automated control strategy. The goal of this development was to create one single reliable algorithm which can be applied to all actuators underneath M1, thus being robust to variations in dynamics, which moreover automatically tunes itself to obtain optimal performance, and tracks the disturbances over time, thereby maintaining its optimality. The basis is a robust linear feedback loop, which is automatically identified within a few seconds. The identified closed-loop system is then used in a real-time estimation of the disturbances, with which an additional performance filter is adapted.^{4,5} This automated controller converges to optimal performance in just a few minutes from the start, and keeps track of changing disturbances. The stability and reliability of the algorithm can easily be manipulated, and has been demonstrated over multiple days.

In this paper the basics of this novel control approach will be explained further. Furthermore, experimental test results will be presented which demonstrate the feasibility of this approach on future segmented primary mirror telescopes. To this end the algorithm has been implemented on a single segment Position Actuator (PACT) prototype⁶ for the E-ELT, and validated in an M1-relevant test setup.⁷ The result will demonstrate convergence to sub-nanometer accuracy, which surpasses earlier results with classic linear controller design, even though this design was explicitly tuned towards the disturbances.⁷ These tests illustrate that with this control approach all 2394 actuators could in principle be automatically tuned to an optimal performance within just a couple of minutes after the press on a single button, and kept optimal even when the wind forces or structural vibrations change over time.

2. METHODOLOGY

This section will give a short overview of the considered E-ELT M1 actuator and its test setup, before describing the suggested control approach in more detail.

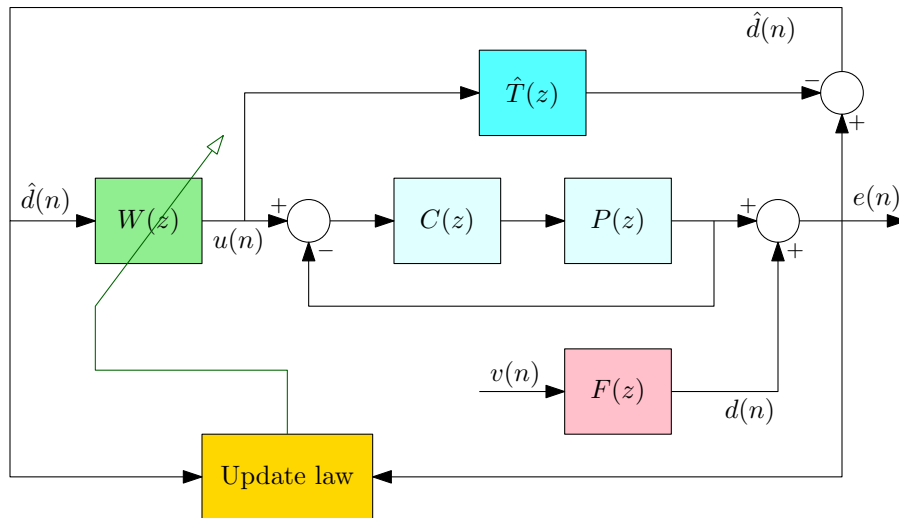


Figure 2: Overview of the proposed control architecture.

2.1 Actuator and test setup

Over the last couple of years TNO has designed and realized various M1 Position Actuator (PACT) prototypes for the E-ELT,^{6,7} a picture of one of them is shown in figure 1(a). As shown in the schematic of figure 1(b), these actuators use a dual-stage concept, where the soft fine stage (a voice coil) acts parallel to an internal spring (meant for gravity compensation) and the coarse stage is needed to achieve a large stroke. PACT is equipped with a nm-resolution on-axis encoder, used as feedback sensor; the aim is to achieve 1.7 nm RMS accuracy on this sensor in closed loop, in the presence of external disturbances. The actuators have been tested in the setup illustrated in figure 1(c); in there PACT pushes against a 90 kg dummy mass via a strut stiffness matching the one in a typical M1 support structure (M1SS). The dynamic response of this assembly is representative for the dynamics of a single servo loop of a mirror segment, and is therefore considered a relevant setup for validating the adaptive control approach presented in section 2.2. Further details on the PACT assembly and its dynamics have been presented during the 2014 SPIE Conference.^{7,8}

2.2 Control architecture

The main goal of the research presented in this paper is to develop a control approach for segmented mirror actuators such as presented in section 2.1, which is applicable to a range of different systems with slight variations in dynamics, while achieving the best possible performance in a changing environment, without taking knowledge of the disturbances into account. Moreover, it should be robust and fail-safe, hence the architecture is set up as a two-stage concept: a robust inner control loop with a standard linear controller (with a fall-back performance of roughly 10 nm accuracy), and an adaptive ‘add-on’ which tracks the time-varying disturbances to further optimize the performance (to sub-nm accuracy).

2.2.1 Building blocks

The essence of the control architecture is illustrated on the basis of figure 2. The open-loop dynamics of the system, i.e. PACT in assembly with the M1 support structure (or with the test setup in this case), is represented by $P(z)$, where z is the discrete-time delay operator. Note that this $P(z)$ is unknown and uncertain, in the sense that it varies spatially along the primary mirror of the telescope. Still, we assume that it is possible to define a basic linear controller $C(z)$ which stabilizes all possible plants. This controller has two purposes:

- it reduces plant variations via the linearizing effect of the closed loop;
- it provide a basic (or fall-back) performance in case the adaptation is turned off.

The performance variable is $e(n)$, which is the output of the plant $P(z)$ plus some unknown disturbance $d(n)$. We can represent this disturbance as stemming from a white noise source $v(n)$ filtered by $F(z)$, which is assumed to be unknown, i.e. knowledge of $F(z)$ is not required in the controller design. Notice that this scheme is also valid for tracking, i.e. when $e(n)$ is the difference between the plant output and a setpoint, since the setpoint signal can be rewritten as an additional disturbance.

We now define a control signal $u(n)$ to be added to the input signal of $C(z)$, so that the true transfer function from $u(n)$ to the performance variable $e(n)$ is given by

$$e(n) = \frac{P(z)C(z)}{1 + P(z)C(z)} \cdot u(n) + F(z)v(n) = T(z)u(n) + F(z)v(n), \quad (1)$$

where $T(z)$ is the true complementary sensitivity of the closed loop. Although $T(z)$ is unknown, we can quickly (in less than half a minute) make a reliable finite impulse response (FIR) estimate $\hat{T}(z)$, by perturbing $u(n)$, monitoring $e(n)$ and making a least-squares FIR fit between this input and output. This $\hat{T}(z)$ is then used to estimate the undisturbed plant output; its difference with the true output then gives an estimate $\hat{d}(n)$ of the actual disturbance in real-time. This approach is known as Internal Model Control (IMC).⁴

If the disturbance model $F(z)$ would be known, one could then define an optimal FIR filter $W(z)$ (also known as a Wiener filter⁴) between $\hat{d}(n)$ and $u(n)$ which will then minimize the power in the performance variable $e(n)$. In our case however, we assume $F(z)$ is unknown and time-varying, hence the filter $W(z)$ needs to be adapted. The cost function J for the adaptation is defined as

$$J = E \{e^2(n)\}, \quad (2)$$

which will be minimized by adaptation of the filter coefficients $\mathbf{w}(n)$. This update law can be defined in many ways, e.g. via a steepest descent method

$$\mathbf{w}(n+1) = \mathbf{w}(n) - \mu \cdot \frac{\partial J}{\partial \mathbf{w}(n)}, \quad (3)$$

where the new filter coefficients depend on the previous ones and the instantaneous estimate of the gradient of the cost function. The latter simply follows from (2) and (3). Note that no additional external perturbation signals are needed to perform the gradient estimation or adaptation; the actual disturbance $d(n)$ is sufficient.

2.2.2 Algorithm reliability

The application of IMC introduces a second closed loop to the control architecture, apart from the basic linear feedback loop $T = \frac{PC}{1+PC}$. This feedback loop involves $W(z)$ and both $T(z)$ and its estimate $\hat{T}(z)$, so its loop gain $L(z)$ is given by

$$L(z) = \left(\hat{T}(z) - T(z)\right) W(z). \quad (4)$$

This loop is guaranteed to be stable if and only if the Nyquist plot of $L(j\omega)$ does not encircle the point $(-1, 0)$. This is clearly the case if $L(z) < 1$ for all frequencies (a result known as the small gain theorem⁹), which is true if there is a perfect fit of the real system $\hat{T}(z) \approx T(z)$ or if $W(z)$ is relatively small. In this case we assume there is no perfect fit, and the real system could somewhat change over time, so to guarantee stability of the loop we need to constrain $W(z)$. We can do so by adding a soft constraint to the cost function so that

$$J = E \{e^2(n)\} + \beta \mathbf{w}^T(n) \mathbf{w}(n). \quad (5)$$

The gradient of the new cost function can again easily be derived, which then introduces a so-called leakage term in the update law (3), i.e. the term $\mathbf{w}(n)$ will be pre-multiplied by a factor $1 - \alpha(\beta)$, with $0 < \alpha(\beta) < 1$ typically relatively small, so that the growth of the filter coefficients is inhibited. Hence *beta* is essentially a parameter to tune the performance and robustness of the algorithm; the smaller β the better the performance, the larger β the higher its reliability.

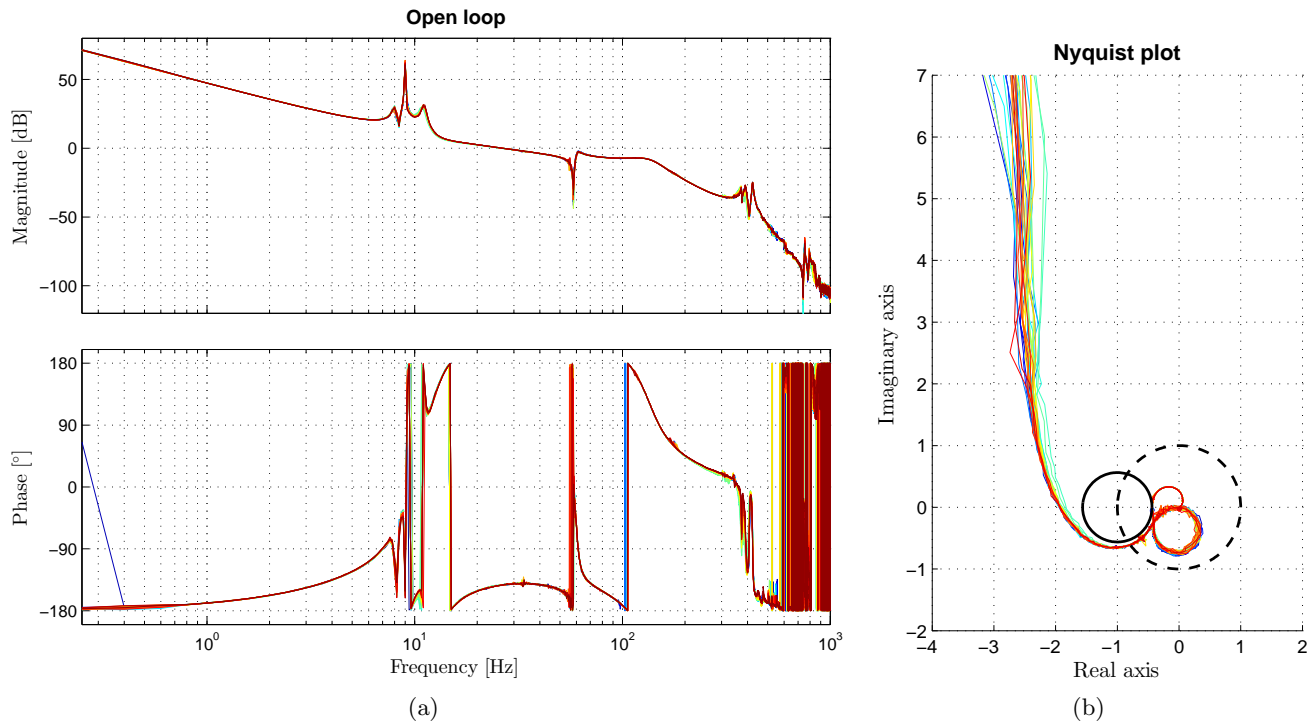


Figure 3: Frequency response of the open loop of the robust inner feedback loop. The various lines are drawn on top of each other represent different measurements of the plant $P(z)$.

2.3 Implementation

The whole algorithm runs on a 5 kHz real-time control platform. It essentially follows three steps:

1. closing the inner feedback via the controller $C(z)$;
2. make the FIR estimate $\hat{T}(z)$ within half a minute;
3. activate the FIR filter $W(z)$ and its adaptation.

The inner controller $C(z)$ is tuned off-line, in such a way that it stabilizes all imaginable plant dynamics with sufficient margins. To fully make use of the linearizing effect of the feedback loop, these margins should mainly be affected by the known controller dynamics instead of the uncertain plant dynamics, so that indeed the closed-loop transfer functions are nearly identical for all imaginable plants. The specific controller design for the results presented in section 3 is illustrated in the Bode and Nyquist diagrams of the open loop $P(z)C(z)$ in figure 3. Compared to high-performance linear controllers for the same hardware,⁷ this $C(z)$ achieves a lower bandwidth (26 Hz instead of 54 Hz), does not have a high gain at the 400 Hz modes, has a better modulus margin (5 dB instead of 6 dB) and yields a very well-defined open loop around 0 dB. Hence, it is much more robust, reduces plant variations, but will also perform less well.

The suggested control approach is meant to minimize the error while disturbances are acting on the system. In our experimental setup there are two different disturbance sources:

Ground vibrations. The setup is located in a normal laboratory, on the second floor in a non vibration-isolated building in an urban area. Vibrations in the ground are transmitted to the encoder readout, despite some level of passive isolation in the system. Ground vibrations are always present, but are typically time-varying; slamming doors, trucks driving by, presence of people, etc., are continuously changing the signature of this disturbance.




Figure 4: Power spectral density (PSD) of the added wind force disturbance.

Wind forces. On a real telescope there are wind forces acting on the segments; in our laboratory we simulate that effect by adding a force disturbance signal in our control loop. This force is obtained by filtering a white noise signal, yielding the spectrum shown in figure 4, which is based on the specified environmental conditions at the E-ELT site.^{7,10} Notice that the force spectrum is time-invariant and mostly low frequent.

Two initial results are shown in figure 5. The blue line depicts the power spectral density (PSD) of a specific open-loop measurement as measured on the internal sensor, including all disturbances but without any feedback, typically summing to about 1 μm RMS movement. The green line represents the error PSD with the inner feedback loop active. It sums to about 10 nm RMS; notice that the integrator present in $C(z)$ already compensates for most of the wind forces. This 10 nm is still far from the 1.7 nm target, but it provides a fall-back performance in the case the additional adaptive loop is turned off, which makes the proposed control architecture more or less fail-safe.

3. RESULTS

In this section the results of the ‘add-on’ adaptive part of the control approach will be presented. This second loop essentially takes the performance of the inner loop as a starting point, and further improves the error automatically.

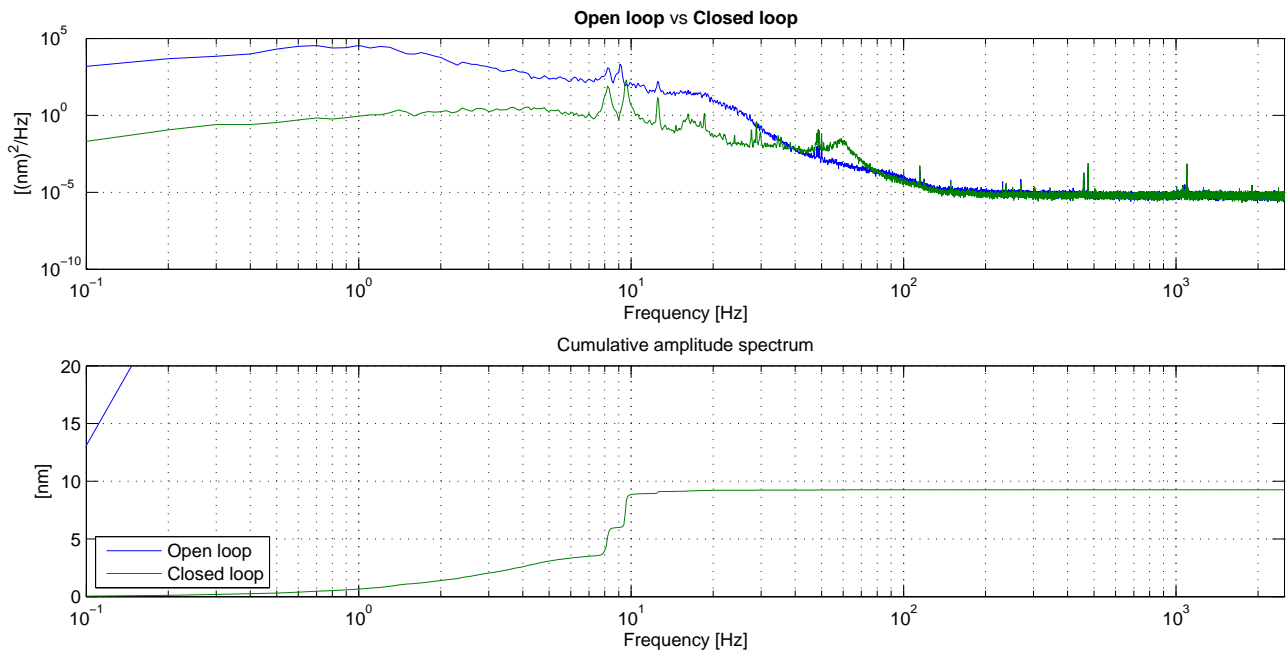
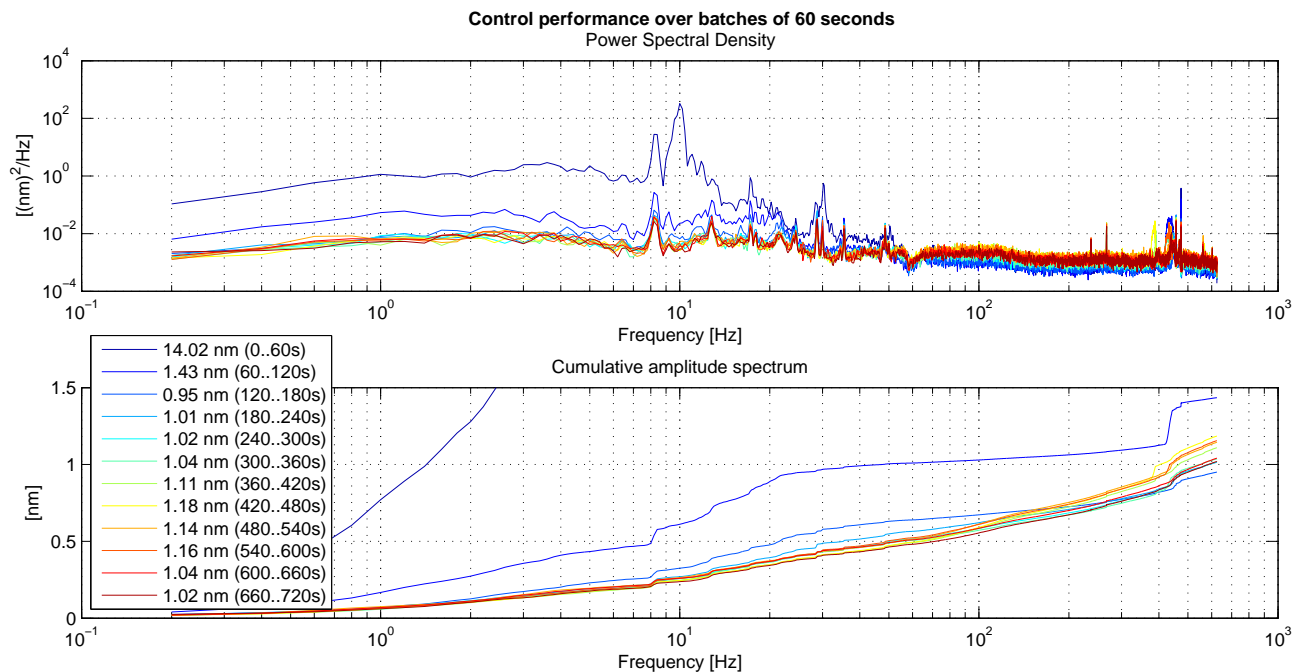
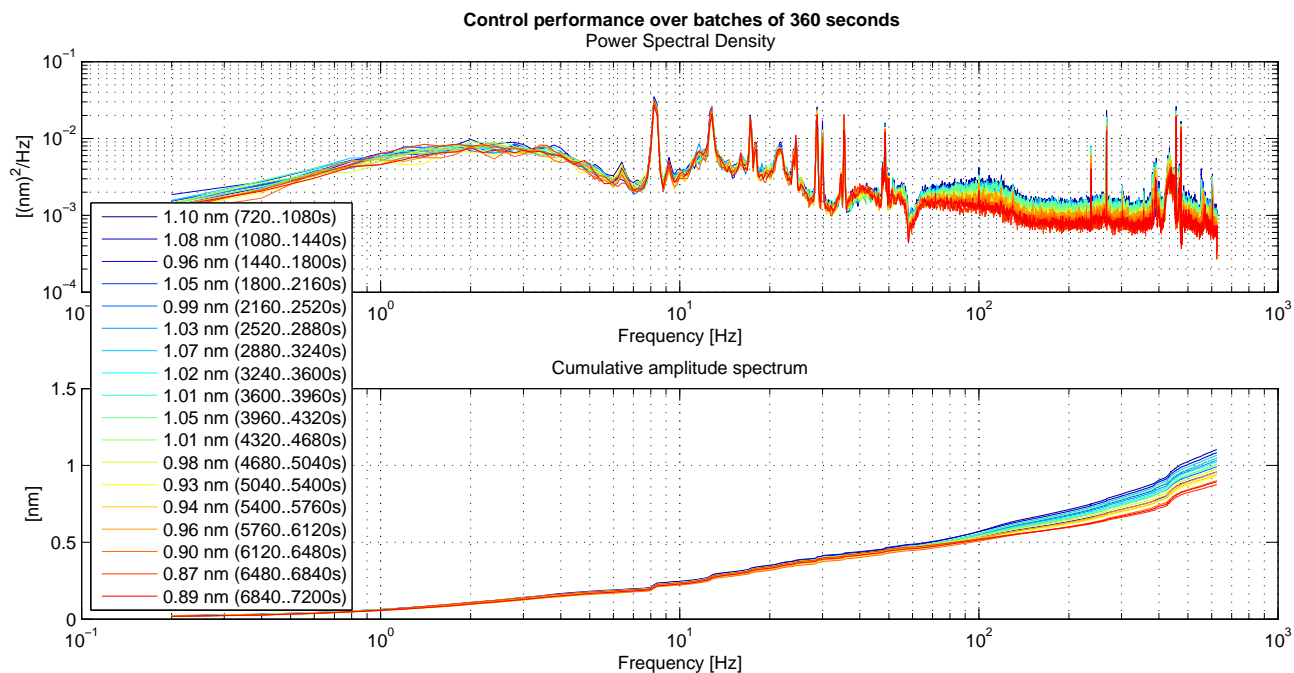


Figure 5: PSD comparison between the open-loop error (blue) and the error with only the inner controller $C(z)$ active (green).



(a) PSDs over 60s during the first 12 min



(b) Subsequent PSDs over 360s up to 2 h

Figure 6: Power spectral densities and their left-integrated RMS of the closed-loop error with the adaptive algorithm turned on. The controller ran at 5 kHz but the output signals have been downsampled with a factor 4 in order to be able to store two hours of data.




Figure 7: Time trace of the closed-loop error during the first and last two minutes of the experiment.

The effect of the adaptive loop right after switch-on is illustrated in figure 6(a). It shows the average PSDs over 60s for the first 12 min after the adaptation is turned on. There is no *a priori* information in the filter $W(z)$, hence all its coefficients are initially set to zero, i.e. $\mathbf{w}(0) = \mathbf{0}$.

In the first minute the performance is slightly less than with only $C(z)$ active, but soon afterwards the error quickly decreases. After three minutes it already converged to about 1 nm. Due to the non-stationary character of the ground vibrations and the further settling of the filter coefficients $\mathbf{w}(n)$, the average error goes up and down a little bit in the first 12 minutes, but never gets larger than 1.2 nm. This is already smaller than the 1.4 nm which had previously been obtained with a carefully tuned high-performance linear controller.⁷

Over a longer time-scale the error improves even more, as demonstrated in figure 6(b). This figure shows average PSD over 360s up to two hours after switch-on of the adaptation. One can see a gradual improvement to just 0.9 nm in this time frame, which nearly twice as good as the target of 1.7 nm. The final PSD still shows some peaks, which most likely originate from ground vibrations, but the left-integrated cumulative spectrum shows that these peaks do not significantly contribute to the total RMS error. Hence, we can conclude that the final PSD is ‘relatively flat’ and thus close to being optimal.

To get an idea of the convergence and end result of the algorithm, figure 7 shows two time traces of the feedback error, both at the beginning and end of the same experiment. These plots confirm that the performance improves dramatically in the first two minutes, and eventually leads to an error of mainly ± 1 encoder resolution bit of 1.2 nm.

Finally, it should be noted that the algorithm has been running reliably for multiple days without encountering any stability issues and keeping the performance below 1 nm. This confirms that the introduction of the leakage α indeed increases the robustness of this adaptive control approach, which makes it applicable to segment actuators in actual telescopes.

4. CONCLUSIONS

In this paper we have proposed a novel control architecture for segmented primary mirror telescopes to automatically achieve optimal co-phasing between the segments. In this architecture a robust inner feedback loop provides a basic fall-back performance for each actuator, which at the same time reduces the variations in the actuator dynamics along the primary mirror. The second ‘add-on’ outer loop then first identifies the actual closed loop, uses this to estimate the disturbance in real-time, and then adapts a second control filter to minimize the control error. This control approach has successfully been applied on an actual segment actuator in an M1-relevant

test setup, demonstrating convergence to sub-nanometer accuracy within just a few minutes. Moreover, the algorithm has been observed to run stably for multiple days, which proves its reliability.

The benefit of this approach is that it provides one single algorithm which is applicable to a range of expected actuator dynamics, and does not need any explicit knowledge of the disturbances. Hence, there is no need for a dedicated controller design for each actuator, and changes in disturbances are tracked automatically. Essentially, for E-ELT this implies that near-optimal performance on all 2394 segment actuators can be achieved with just the press of a single button.

REFERENCES

- [1] Jared, R. C., Arthur, A. A., Andreae, S., Biocca, A. K., Cohen, R. W., Fuertes, J. M., Franck, J., Gabor, G., Llacer, J., Mast, T. S., Meng, J. D., Merrick, T. L., Minor, R. H., Nelson, J. T., Orayani, M., Salz, P., Schaefer, B. A., and Witebsky, C., “The W.M. Keck telescope segmented primary mirror active control system,” *Proc. SPIE* **1236**, 996–1008 (1990).
- [2] Erm, T., “Technical specification for the manufacturing and delivery of prototype position actuators for the primary mirror of the European Extremely Large Telescope (E-ELT),” Tech. Rep. E-SPE-ESO-142-0011, issue 2, ESO (2009).
- [3] Thompson, P. M., MacMynowski, D. G., and Sirota, M. J., “Control analysis of the TMT primary segment assembly,” *Proc. SPIE* **7012**, 70121N (2008).
- [4] Elliot, S. J., [*Signal Processing for Active Control*], Academic Press, London, UK (2001).
- [5] Berkhoff, A. P., “A filter apparatus for actively reducing noise.” Patent WO 2007/064203 A1 (7 June 2007).
- [6] Kamphues, F., Nijenhuis, J. R., den Breeje, R., van den Dool, T. C., and Ponsioen, J., “PACT: the actuator to support the primary mirror of the ELT,” *Proc. SPIE* **7018**, 70180Z (2008).
- [7] Witvoet, G., den Breeje, R., Nijenhuis, J., Hazelebach, R., and Doelman, N., “Dynamic analysis and control of mirror segment actuators for the european extremely large telescope,” *Journal of Astronomical Telescopes, Instruments, and Systems* **1**(1), 019003 (2015).
- [8] Witvoet, G., den Breeje, R., Nijenhuis, J., Hazelebach, R., and Doelman, N., “High performance control of mirror segment actuators for the European Extremely Large Telescope,” *Proc. SPIE* **9145**, 91451S (2014).
- [9] Skogestad, S. and Postlethwaite, I., [*Multivariable Feedback Control: analysis and design*], John Wiley, Chichester, UK (2005).
- [10] Koehler, B., “E-ELT environmental conditions,” Tech. Rep. E-SPE-ESO-313-0111, issue 5, ESO (2012).