

Distributed Force Control of Deformable Mirrors

Douglas G. MacMynowski¹, Rikard Heimsten^{2,*}, Torben Andersen²

¹ Control and Dynamical Systems, California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91125, USA;

² Lund Observatory, Lund University, SE-221 00 Lund, Sweden

Large (>1 m) deformable mirrors are attractive for adaptive optics on ground-based telescopes; the mirrors typically have hundreds or thousands of actuators. The use of force actuators instead of position actuators has the potential to significantly reduce total system cost. However, the use of force actuators results in many lightly-damped structural resonances within the desired bandwidth of the control system. We present a robust control approach for this problem and demonstrate its performance in simulation. First, we demonstrate that high-bandwidth active damping using velocity feedback from mirror sensors that are not quite collocated with the actuators can be robustly implemented, because at sufficiently high frequencies the structural dynamics enter an “acoustic” limit, where the half power bandwidth of a mode exceeds the modal spacing. This is important, because the system can be made less expensive using sensors placed in between actuators rather than collocated with each actuator. Introduction of active damping leads to a much easier problem for subsequent position control. It is known that a position control system in which each of the actuators is controlled using feedback from a collocated sensor can be made robustly stable. However, the resulting performance at high spatial frequencies is poor because there is no shared information between neighbouring actuators. In contrast, global control gives excellent performance but lacks robustness to model uncertainty. We introduce an innovative local control approach, which significantly improves the high spatial frequency performance without the robustness challenges associated with a global control approach. The overall approach is demonstrated

to provide excellent command response suitable for an adaptive optics outer loop.

Keywords: Adaptive optics, deformable mirror, distributed control, active damping, local control

1. Introduction

Most modern ground-based telescopes have adaptive optics (AO) for compensation of atmospheric blurring or telescope aberrations. Adaptive optics encompass one or more deformable mirrors (DMs), each with many actuators, to compensate for distortion of the wavefront of the incoming light. So far, most DMs have been small, typically with a diameter of tens of centimeters or less. However, large DMs with a size of 1–4 m are highly attractive for the new generation of Extremely Large Telescopes, because the DMs can then be integrated directly into the optical telescope, avoiding use of post-focus relay optics. The present paper outlines an approach for control of such a large DM.

DMs up to about 1 m exist and larger DMs are planned. Stiff “position” actuators of piezo-electric, electrostrictive or magnetostrictive type have so far successfully been used to deform a DM with a high temporal bandwidth. Because the DM is effectively constrained at the actuator locations, there are no significant problems related to dynamic performance. On the other hand, the use of such actuators is costly due to the tight fabrication tolerances involved. In

*Correspondence to: R. Heimsten, E-mail: rikard@astro.lu.se

contrast, a DM with soft force actuators is more difficult to control because the dynamics of the mirror structure enter directly into the control loops. However, DMs with force actuators have the potential for being inexpensive because the mechanical tolerances can be more relaxed. Since the cost of DMs may be dramatic, we have studied approaches for use of low-cost force actuators.

An Italian group has played an important role in the development of large DMs with force actuators, and has set up a control strategy [3, 13] for their systems. Electronic damping is established by differentiating the position feedback from a collocated sensor at the location of the actuator. In addition, a combination of feedforward and position feedback is used to track commands from an external AO control loop. The mirror has a thickness of 1–2 mm and is made of polished Zerodur. A British group has studied an alternative design using composites [5] but has not published a detailed control strategy for force actuators.

For development of a low-cost DM, we have studied a system using force actuators attached to the back of the mirror through suction cups and non-collocated position sensors using electret microphones in rubber bellows touching the back of the mirror [1] much like stethoscopes. The DM itself can also be inexpensively constructed using slumping techniques [1]. Due to the use of inexpensive components and low tolerances, such a system has the potential for being much less expensive than corresponding systems built or proposed until now. However, the system is more complex from a controls point of view because the structural dynamics of the mirror are part of the control loops. Also, to provide sufficient space, the force actuators and position sensors are not collocated as will be shown later. The focus of this paper is on demonstrating a control approach for this problem that gives sufficient command-response bandwidth for use within an AO system.

The control approach in [3, 13] uses single-input-single-output (SISO) feedback from mirror deflection sensors that are collocated with the actuators (“collocated control”). While guaranteed to be robust, feedforward is required to obtain adequate performance at high spatial frequencies, requiring careful calibration. Miller and Grocott [11] present a global control approach for a similar flexible DM, combining several innovations, including the use of a full state-space controller and a circulant approach to reduce the resulting computations. While overcoming the limitations of collocated control, the challenges with this model-based approach are likely to come from robustness; any control that relies on some particular model information has the potential for insufficient robustness to uncertainty in that information. A distributed control approach has also been simulated for a large deformable mirror [6, 4]. This approach leads to a low-computation decentralized approximation to the global optimal control solution, but since it approximates the system as infinite

in extent, it does not correctly deal with the boundaries, which can limit the achievable performance.

The approach taken here is to use the minimum information required about the structural dynamics to ensure robustness to modeling uncertainty, while incorporating sufficient information to provide performance, and without relying on accurate calibration. The approach can be divided into two steps. First, active damping using velocity feedback effectively compensates for the lightly-damped structural dynamics. Second, a position control system is added. Due to the active damping, the position control system can be based upon local controllers, in turn providing excellent command response suitable for an outer adaptive optics feedback loop based on wavefront information. Hence, in this paper we present two key innovations. First we show that active damping can be implemented using sensors that are not perfectly collocated, and we compute an explicit bound on the maximum spatial separation that can be tolerated between actuators and sensors, using an acoustic argument. This allows the use of realistic and inexpensive hardware. And second, we present a robust local position control strategy overcoming the high-spatial frequency limitations of collocated position control without introducing the robustness concerns of a model-dependent global feedback strategy; this is similar in spirit to the local estimation strategy in [10].

This control approach is also of potential interest for future space applications. For weight reasons, large spacecraft mirrors must necessarily be thin, calling for a control system to set the correct shape of the mirror during operation in the presence of significant structural dynamics within the control bandwidth.

The next section outlines the control problem in general and a simulation example used throughout to illustrate the approach. The velocity and position feedback concepts are discussed in Sections 3 and 4 respectively.

2. Control Problem

2.1. Problem Characteristics

The objective of the control strategy developed herein is to enable a low-cost concept for a large deformable mirror that uses force actuators attached to the back of the mirror with suction cups. Use of low-cost force actuators in combination with a thin mirror leads to a lightly damped system with structural resonance frequencies within the desired control bandwidth. Feedback from the wavefront sensor in the external adaptive optics loop will not be sufficient to confront the problem with structural resonances. Hence, an additional feedback loop using sensors on the back of the mirror is needed as shown in Fig. 1; the focus of this paper is the design strategy of this feedback loop.

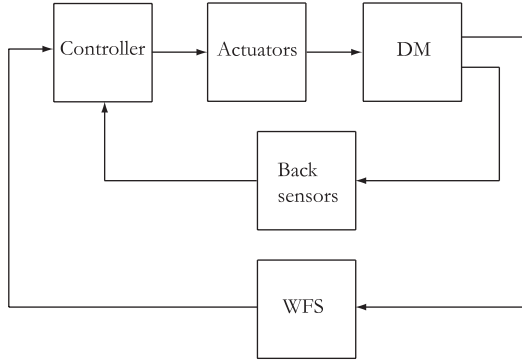


Fig. 1. Schematic of overall control architecture, including inner control loop of a flexible, deformable mirror, and outer adaptive optics control loop using wavefront sensor feedback.

We plan to use electret microphones in bellows at the back of the mirror to sense mirror deflection. For practical reasons, it is not attractive to have the actuators and back sensors precisely collocated. One of the contributions herein is to demonstrate that robust performance does not require this. The impact of finite actuator/sensor bandwidth and delays from discrete-time implementation can be understood in terms of their impact on phase margin. Other practical details regarding actuator/sensor design are discussed in [1]. The prototype actuator in [1] has sufficient force to give a low-spatial frequency amplitude of about 1 mm for the mirror parameters used in the simulations here, and a maximum difference in excursion between two adjacent actuators of about $1.5 \mu\text{m}$; this is more than enough for AO.

The requirements on the control algorithm for the DM are determined by the characteristics of the AO system, and are discussed in [1]. The principal requirement is the temporal command bandwidth. In order to allow a typical closed-loop AO bandwidth of roughly 50 Hz, the DM must respond to actuator commands with relatively flat frequency response and at most have a $20\text{--}30^\circ$ phase lag at that frequency. (Note that the AO bandwidth is typically limited by measurement (photon) noise, and not by phase lag.) This requires that the system bandwidth for the DM be higher than the desired AO bandwidth. Herein we evaluate command response accuracy by assuming all spatial frequencies correctable by the DM are equally weighted. The range of correctable spatial frequencies is determined by the number of actuators.

2.2. Simulation Example

The example used throughout this study is a 2 mm thick, 1 m diameter deformable mirror fixed in its center and made of the material borosilicate. Material properties are given in Table 1. The actuator and back sensor locations are shown in Fig. 2. A finite element model of the mirror

Table 1. Parameter definitions and values.

Parameter	Definition	Value
E	Young's modulus	$63 \times 10^9 \text{ Pa}$
ρ	Density	$2.23 \times 10^3 \text{ kg/m}^3$
ν	Poisson ratio	0.2
h	DM thickness	2 mm
I	Moment of inertia	$h^3/12$
D	Bending stiffness	$EI/(1 - \nu^2)$
A	DM area	$\pi(0.5 \text{ m})^2$
ζ	Damping ratio	1%

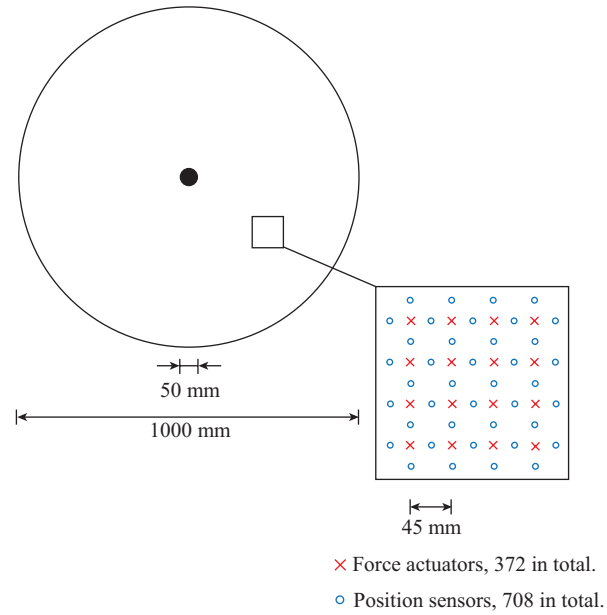


Fig. 2. Topology of distribution of force actuators and position sensors over the back of the 1 m diameter deformable mirror used in the examples.

has been set up using the software package “Comsol Multiphysics” approximating the mirror with a plane shell. The mesh includes about 5000 nodes, of which 1080 are placed at the actuator and back sensor locations. The remaining nodes were automatically generated with the constraint of a maximum node distance of 1.8 cm. Each node has six degrees of freedom (DoF), i.e. translation along three mutually perpendicular axes and rotation around the same axes. The dynamic behavior of the structure is described by the differential equation:

$$\mathbf{M} \frac{d^2 \delta}{dt^2} + \mathbf{D} \frac{d \delta}{dt} + \mathbf{K} \delta = \mathbf{f}$$

where \mathbf{M} , \mathbf{D} , \mathbf{K} are the mass, damping and stiffness matrices, \mathbf{f} the force vector, and δ is the vector of angular and translation displacements. The lowest eigenfrequency of the structure is 5.6 Hz with another 658 eigenfrequencies almost evenly distributed up to 5000 Hz (consistent with

Eq. (3)). In the general case, there will be four sensors located at the same distance from a given actuator, see Fig. 2. For the purpose of position feedback for the actuator, we apply an average value of the reading of the four sensors. In the following, we will refer to the average of the four sensors, y , as the “sensor” signal for the actuator. Use of this non-collocated sensor and actuator scheme leads to a phase lag between the actual movement of an actuator and the signal from a nearby sensor; this can be estimated from the wave speed.

The full model of the mirror with nearly 30000 DoFs is computationally impractical, so model reduction has been performed as follows:

- First, Guyan reduction was utilized to reduce the number of DoFs to three for each node, retaining out-of-plane translation and the two rotations about the in-plane axes. This reduction is applied to the structure matrices \mathbf{K} and \mathbf{M} .
- Next, a further model reduction was performed using modal truncation, removing modes with eigenfrequencies above 5000 Hz. Mode acceleration was applied to include the static contribution from the modes that were omitted by the truncation. The final model is formulated in a structural modal basis.

3. Velocity Control

The purpose of conceptually dividing the feedback into rate (“electronic” or “active” damping) and position control loops is that the latter is more straightforward to design once the former loops are closed. The goal of the rate feedback is to compensate for the lack of damping in the mirror structure. Active damping with collocated and ideal actuators and sensors is straightforward; we here demonstrate that active damping is also possible without perfect collocation, and we quantify the allowable separation. To do so, it is first useful to explore several characteristics of the flexible deformable mirror dynamics. Section 3.1 describes the characteristics of the structural dynamics that enable this velocity control approach and Section 3.2 outlines the rate feedback.

3.1. Wave and Impedance Considerations

For systems with relatively few modes, a modal description is useful. However, with many modes within the control bandwidth, useful intuition can also be obtained with a wave-based description of the structural dynamics. Resonances arise due to constructive interference of waves reflecting off system boundaries. At high frequencies, small variations in the material properties, interaction with actuators, or boundary details will result in significant

uncertainty in the structural resonances, making any approach that is reliant on exact knowledge of the resonances non-robust. However, the relevant properties of the mirror near an actuator depend only on the local structure properties, and are much less sensitive to modeling errors.

For the purpose of understanding the dynamics, it is sufficient to model the deformable mirror by an ideal plate. The bending wave (group) speed in the plate at angular temporal frequency ω is

$$c = 2\sqrt[4]{\frac{D}{\rho h}}\omega^2 \quad (1)$$

where parameter definitions are given in Table 1. For the properties used in the simulation example, $c = (3.6 \text{ ms}^{-1/2})\sqrt{\omega}$.

The drive-point mobility, i.e. the transfer function from force to velocity at the same location of an infinite plate is [2]:

$$G = \frac{1}{8\sqrt{D\rho h}} \quad (2)$$

This is also the “average” or dereverberated transfer function of a finite plate [9] as evident in the representative transfer function shown in Fig. 3. If the structural damping is sufficiently high, then disturbances propagating towards the boundary of the mirror dissipate. There is no significant reflection, and therefore, except at very low frequencies, the structure response is the same as that of an infinite plate.

A critical feature also evident in Fig. 3 is that with any non-zero damping, then at sufficiently high frequency, the deviation between the response of the lightly damped system and that of the dereverberated or infinite system decreases. The response is no longer dominated by individual modes, but rather at any frequency, there are sufficiently many different modes participating that the transfer function is no longer characterized by sharp resonant and anti-resonant peaks in either the magnitude or phase. The half-power bandwidth of each mode is $2\zeta f_m$, where ζ is the modal damping and f_m the undamped eigenfrequency for the mode. The “acoustic” behavior occurs for frequencies where the half-power bandwidth of a mode exceeds the average modal spacing by a factor of two or three (i.e. 2–3 modes excited at any given frequency). For a plate with area A , the average modal spacing [8] is:

$$\Delta f = \frac{2}{A}\sqrt{\frac{D}{\rho h}} \quad (3)$$

which is roughly 8 Hz for the simulation parameters here. Thus the transition to acoustic behavior begins around

$$f_a \simeq 2\sqrt{\frac{D}{\rho h}} \frac{1}{A\zeta} \quad (4)$$




Fig. 3. Representative transfer function from force to position at one actuator for a 1 m mirror, compared to the theoretical infinite plate response (dashed line). See Section 2.2 and Table 1 for simulation parameters. The high frequency decrease in phase relative to the theoretical response is due to sensor non-collocation.

which for the parameters here is ~ 800 Hz; above this frequency we should expect to see a smoother transfer function arise (c.f. Fig. 3).

3.2. Active Damping

It is well known that damping can be added using collocated rate feedback (e.g. [12]):

$$u_v = -K_v \dot{y}_a \tag{5}$$

where y_a is the deflection at a given location, K_v is a gain, and u_v is the command to a force actuator collocated with the deflection measurement. The optimal gain is given by the inverse of the dereverberated drive point mobility (the transfer function from force to velocity if the system were completely damped with all “reverberations” removed); this follows from impedance-matching arguments (e.g. [9]). The dereverberated mobility for a plate is the same as the mobility of an infinite plate (Eq. (2)). This feedback is thus positive real and hence stability is guaranteed. However, because both the dereverberated system dynamics and feedback are constant with frequency, the loop transfer function will be of order one at all frequencies. Phase lag due to actuator-sensor non-collocation, actuator/sensor dynamics, or delays from the electronic implementation will result in any real system not being positive real at some sufficiently high frequency. The system will still be stable if there is sufficient damping, see e.g. [12, Fig. 5.15]; for the DM here, the trade-off between

damping and actuator-sensor spacing can be explicitly calculated.

Stability can still be guaranteed if the frequency at which positivity breaks down exceeds the acoustic limit described earlier, where the system transfer function is no longer dominated by lightly damped resonances, and the average phase of the force to velocity transfer function is 0° , rather than alternating between $\pm 90^\circ$. The reduced phase variability allows roll-off to be introduced into the rate feedback without stability problems so that the feedback is gain-stable at high frequencies; the reduced gain variability also means that relatively little roll-off (and accompanying phase lag) is required for high-frequency gain-stability.

Stable rate feedback can thus be implemented provided there is minimal phase lag below this frequency. This requires a minimum rate feedback bandwidth of the order of the acoustic limit defined in Eq. (4). Note that this limit frequency decreases as the deformable mirror increases in size, hence making larger mirrors easier to control. Non-collocation will introduce phase lag of 45° at an actuator-sensor spacing of d at frequency $f = c/(8d)$, and thus the maximum actuator-sensor distance is of order $d \leq c/(8f_a)$ where c and f_a are from Eqs (1) and (4) respectively. With the parameters used here, this means that if the spacing between the actuators and sensors is no more than about 22 mm (this is the value chosen here; see Fig. 2), then the phase lost to non-collocation is tolerable at frequency f_a . Any additional phase lag due to actuator/sensor bandwidth and implementation delays need to be accounted




Fig. 4. Performance of rate feedback loop at a representative actuator location (transfer function magnitude from actuator force to sensor position) for (a) open loop, (b) only the collocated loop closed and (c) all loops closed.

for and will reduce the maximum tolerable spacing. Combined with the knowledge that the plant dynamics are relatively smooth above this frequency, which allows roll-off to be introduced, then this demonstrates that robust rate feedback can be used even with non-collocated (and hence inexpensive) sensing. This is a critical enabler for robust control of a distributed force actuated DM.

The performance of the rate feedback loop is shown in Fig. 4 for the simulation example described in Section 2.2. A rate feedback gain in Eq. (5) of 75 is used, slightly lower than the optimal value of $G^{-1} \simeq 100$ (for added robustness). For our simulation, a low-pass filter is included at 1.5 kHz. For a real implementation this might not be necessary because actuator or sensor dynamics will introduce filtering; this frequency would be the minimum actuator bandwidth for the simulation parameters used herein. The three diagrams in Fig. 4 plot the response from force at one actuator location to the corresponding sensor signal y derived from the average of four neighbor sensors. Figure 4(a) shows the open-loop transfer function (same as Fig. 3), which is highly resonant at low frequencies. Figure 4(b) is the transfer function when the rate feedback loop is closed only for the actuator observed. Modes

with non-zero residuals at that location are damped, but of course there is no effect on the anti-resonances, where there is no displacement at that location. Figure 4(c) is a plot of the response when all rate feedback loops are closed. The resonant behavior has been suppressed by adding the active damping and the mirror behaves as if it were infinite: disturbances are dissipated before they propagate across the mirror and reflect back to constructively interfere. This provides a more straightforward system for the design of the position control loop.

4. Position Control

4.1. Collocated Feedback

The purpose of the position control loop is to provide a good command response for the outer adaptive optics loop. In particular, minimal phase lag and a relatively flat frequency response up to at least 50 Hz are essential. A block diagram describing the overall feedback control architecture is shown in Fig. 5. The open-loop plant dynamics are given by \mathbf{G} ; the other blocks will be described below.

The rate feedback, \mathbf{K}_V , described in Section 3 provides electronic damping using SISO feedback for each actuator. Similarly, the simplest position control is collocated SISO feedback. This can be interpreted as implementing electronic springs, and positivity can again be used to guarantee robust stability. This is the approach taken in [13], and corresponds to choosing $\mathbf{Q} = \mathbf{I}$ in Fig. 5. The performance with this approach is shown in Fig. 6, with \mathbf{K}_P from Fig. 5 representing P-controllers with the highest possible gain while maintaining a stable system (a PI control could also be used). The command response at the representative actuator is -8 dB, with the response at the adjacent actuator only slightly smaller. There is substantial coupling in the plate, and while one actuator pushes up to achieve the desired response, the adjacent actuators are all counteracting this to achieve the different response desired at those locations. The net result is that at the maximum

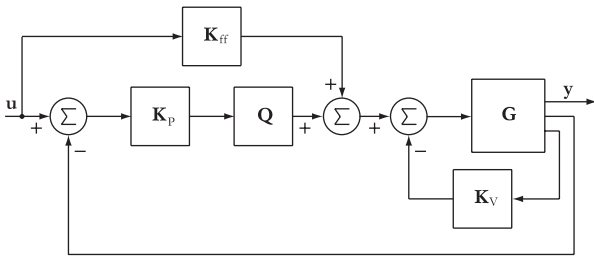


Fig. 5. Block diagram for overall control architecture. See text for definitions.

stable gain, the feedback control has very little effect on this high spatial frequency command response.

The response of the mirror to position commands could also have been plotted for commands at different spatial frequencies, rather than for a command at an individual actuator. The mirror compliance \mathbf{C} is much higher at low spatial frequencies than at high; for our simulation example, the condition number of the static plant response ($\mathbf{G}(0) = \mathbf{C}$) is of the order of 10^5 . This means that there is much higher loop gain for low spatial frequencies, and quite good command response, but low loop gain and hence poor command response at high spatial frequencies.

A possible solution to this is to implement feedforward control (as used in [13]). Based on open-loop calibration of the deformable mirror, the best estimate of the force distribution required to achieve the desired command response is fed to the actuators, and feedback is only required to compensate for errors in the feedforward matrix. The matrix \mathbf{K}_{ff} in Fig. 5 is ideally the stiffness matrix evaluated at the actuator locations. While using a priori information about the system is always beneficial, the low feedback gain at high spatial frequencies means that the deformable mirror is effectively operating in open loop to high spatial frequency commands; this then requires accurate calibration to ensure adequate performance. Note that it is the inverse of the plant that is required for the feedforward control; higher accuracy is required for the plant model than will be obtained on the estimate of the inverse unless an approach is used to




Fig. 6. Command response at representative actuator using collocated position control (solid line), and cross-talk to a neighbour actuator (dashed line) and a distant actuator (dotted line).

measure the inverse directly. Section 4.3 compares the performance robustness of this algorithm with the approach developed below.

4.2. Local Feedback

Clearly if the matrix \mathbf{Q} in Fig. 5 were selected equal to the inverse of the plant at least at zero frequency, the system would be much better conditioned. However, this not only requires the same information as the feedforward solution, but would now be within the control loop, leading to potential instability for small errors in the plant model. However, recall that the problem occurs due to each actuator counteracting the control at the adjacent actuators. The coupling can therefore be significantly reduced using only local information. This enhances robustness relative to using global model information in the feedback. If the control does not require knowledge of the global interaction effects, then it will automatically be robust to uncertainty in these effects; this is verified in Section 4.3.

The goal in choosing the matrix \mathbf{Q} can be stated in one of two (equivalent) ways:

- Given a particular error pattern \mathbf{e} , say $e_k = 1$ at the k^{th} actuator and $e_j = 0, j \neq k$, then rather than only applying a force to the k^{th} actuator (which will produce non-zero response across the entire mirror, see Fig. 7(a)), we wish to apply a force distribution that yields a response pattern to counteract the error, see Fig. 7(c). However, we use only actuators close to actuator k , and will tolerate some error in replicating the error pattern.
- We would like to obtain a matrix \mathbf{Q} that is an approximate inverse to the compliance \mathbf{C} using only local information.

Mathematically, the above problem can be described as choosing the applied force pattern \mathbf{f}_k to minimize the cost function

$$\mathbf{J} = \|\mathbf{C}\mathbf{f}_k - \mathbf{e}\|_2 \quad (6)$$

subject to the constraint that elements of the vector \mathbf{f}_k not in some set Ω_k must be zero. The set Ω_k can be constructed by including all actuators within some specified distance of k . The constrained least-squares problem is equivalent to solving an unconstrained problem with a truncated matrix $\mathbf{C}_{:, \Omega_k}$, indicating that only the columns associated with actuators in Ω_k are retained. Then the optimal $\mathbf{Q}_{\Omega_k, k}$ (giving the appropriate local force distribution to compensate for an error at location k) is the k^{th} row of the pseudo-inverse of the truncated compliance. This is motivated by the local approach used in [10] to develop computationally efficient sparse reconstructor matrices for adaptive optics estimation. A hierarchic or multigrid approach could also

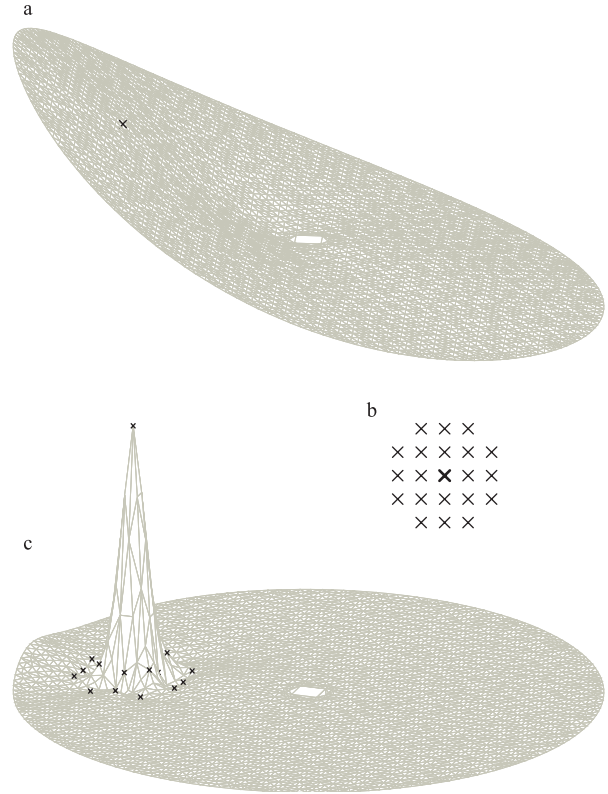


Fig. 7. Overview of the local feedback performance, (a) mirror shape due to a single actuator command, (b) local family group of 21 actuators, (c) mirror shape resulting from unit command to a local actuator family.

be used as in [10, 7] to improve the approximate inverse at larger spatial scales, however, for this application, the larger spatial scales are already adequately controlled.

Recall that the challenge that motivated the need for \mathbf{Q} was that if a single actuator is used to respond to an error, then there is a large global response. Solving a least squares problem where the minimization is over the entire mirror ensures that the response to an error at location k is mostly local. The purpose of the additional actuators within region Ω_k can be interpreted as minimizing the energy propagation away from the local region, as any such energy will dominate the least-squares performance metric.

To illustrate the effectiveness of this approach, we use the set Ω illustrated in Fig. 7(b). The first “ring” of actuators around the center will push in the opposite direction of the center actuator, to counteract the response that would otherwise occur. The second ring of actuators is useful to minimize energy propagation away from the local region. The closed-loop response pattern obtained using this local \mathbf{Q} and the architecture in Fig. 5 is shown in Fig. 9, both for an actuator away from any boundary, and for an actuator location on the boundary. Note that because the additional actuators are required only to minimize energy propagation out of the local region,




Fig. 8. Open-loop response for position-control loop (from an input to \mathbf{Q} to position) at a representative actuator using local control approach. The \mathbf{Q} matrix derived at DC-performance works reasonably well up to about 1000 Hz.

boundary regions where there are fewer actuators in Ω actually perform better.

There is clearly significant potential for exploring different possible sizes for the set Ω ; better performance will occur with larger sets, at the expense of robustness due to the requirement for additional structural dynamics information. The appropriate trade-off for a given mirror and actuator/sensor layout can be obtained by looking at the residual least-squares performance metric from (6) as a function of the size and pattern Ω , and choosing the set beyond which there is diminishing benefit from using additional actuators. Simulations in the next section illustrate that the set chosen here is sufficient.

Referring back to Fig. 5, the alternate interpretation of the purpose of \mathbf{Q} is to improve the conditioning of the position-control plant at zero frequency, $\mathbf{C}\mathbf{Q}$. Using the local set of actuators in Fig. 7(b), the condition number of $\mathbf{C}\mathbf{Q}$ is reduced from $\sim 10^5$ with $\mathbf{Q} = \mathbf{I}$ to ~ 2 , a remarkable improvement using at most 6% of the actuators to construct each column of the approximate inverse.

Using the local solution \mathbf{Q} , the response from the input of the position-control plant $\mathbf{G}\mathbf{Q}$ to the output is now much closer to the identity matrix; the output response distribution is similar to the input distribution. A representative open-loop transfer function for high spatial frequency for the plant is shown in Fig. 8; it is straightforward to design a stable SISO integral control loop for each degree of freedom. The final result, shown in Fig. 9, indicates a flat response up to 200 Hz for actuators at the edges and up to 100 Hz for the remaining actuators. The step response of the mirror at every actuator location to a unit command at

a representative actuator is shown in Fig. 10. It is evident that there exists little cross-talk between a commanded actuator and the actuators outside the local family.

Increasing the integral gain leads to instability due to cross-talk between local families and local dynamics within each local region Ω . The command response and cross-talk are both excellent for use with an outer AO loop, with roughly 25° phase lag at the expected AO bandwidth of 50 Hz.

4.3. Robustness

The robustness of the control system depends on the accuracy of \mathbf{Q} , or equivalently, on the uncertainty in the plant. The matrix \mathbf{Q} can be determined either by modeling or experimentally through poking of actuators. Both options will lead to somewhat incorrect elements of \mathbf{Q} .

For a lightly damped structure, small uncertainties in the resonant frequencies will lead to large uncertainty in the transfer function; this type of structured uncertainty in the system is poorly represented by unstructured uncertainty models. We instead evaluate stability- and performance-robustness for this system by explicitly evaluating the effect of multiplicative uncertainty in the resonant frequencies, while keeping the damping ratio constant.

The metric we use to compare the command-response performance of different algorithms in the presence of model uncertainty is to compute the mean-square error over all spatial degrees of freedom and within the 0–50 Hz expected bandwidth of the AO system. Fig. 11 compares




Fig. 9. Command response using local control approach at two different actuator locations away from (a) and at outer boundary (b), illustrating collocated response (solid line), and cross-talk to a neighbouring (dashed line) and a distant actuator (dotted line). The half-power bandwidth exceeds 100 Hz, and the phase lag at 50 Hz is less than 25° .

this error index in the presence of bounded uncertainty in the structural resonances for three different control approaches:

1. A guaranteed-stable collocated position control (the matrix \mathbf{Q} is the identity matrix). Acceptable performance requires feedforward, and the fully-populated
2. The local control approach presented herein (the matrix \mathbf{Q} is sparse and local). While not required for acceptable performance, the same feedforward is used

best model-based estimate of the required force distribution pattern is used, constructed from the nominal model without uncertainty; this is similar to the approach used in [3, 13].




Fig. 10. Step response for the local approach suggested here, with feedforward. The response for a commanded actuator (thick line), the corresponding response of the local family (dashed lines) and the cross-talk to actuators outside the local family (thin lines).

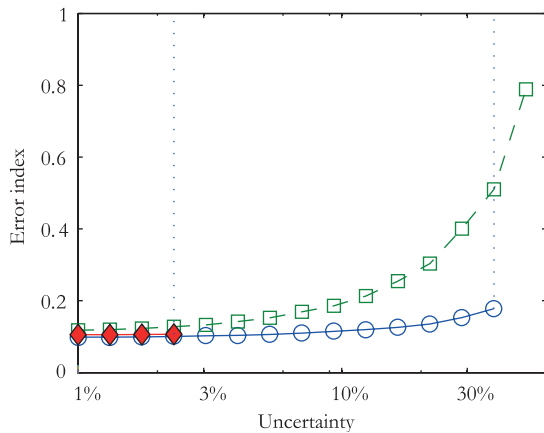


Fig. 11. Performance robustness comparison for three different controllers: the local approach suggested here (“○”), collocated control (“□”, dashed), and with global feedback of all sensors to all actuators (“◆”). All are plotted assuming feedforward information is used. Uncertainty is bounded random errors in the mirror resonant frequencies, and the error index is the mean-square command-response error averaged over all controlled spatial degrees of freedom and temporal frequency from 0 to 50 Hz. The error index is infinite (indicated by the dashed line) if the closed-loop system is not stable, which occurs for both the local and global feedback.

Q1

as above in order to make a direct comparison between the approaches.

3. A global control approach, where **Q** is the best model-based estimate of the required force distribution, again for the nominal model. Again, feedforward information is also assumed.

The three approaches above are listed in order of how much information is needed in the feedback path about the system dynamics. Other model-based control algorithms are certainly possible that would be more robust than the global approach we compare to here. The vertical dashed

lines in Fig. 11 indicate the uncertainty at which the local and global approaches are not guaranteed to be stable. The local position control approach proposed here gives nominal performance close to that of the global approach, and maintains stability for 40% uncertainty in resonant frequencies, more than 10 times what can be tolerated with this global approach. Because there is sufficient temporal feedback bandwidth on almost all of the spatial degrees of freedom controlled by the DM, the performance does not degrade as quickly with increasing model error as for the collocated approach, where performance on many high spatial-frequency degrees of freedom is essentially obtained in open-loop with only feedforward correction.

Note that increasing the size of the actuator set (number of actuators used in each local family) would improve performance at the expense of robustness (as it requires more model information). Since the performance with this local set is nearly the same as with full global information, there appears to be no advantage to increasing the set size.

5. Conclusions

Traditional actuators for large deformable mirrors call for tight construction tolerances, which imply high costs. We have here demonstrated a stable and robust control strategy that enables a low-cost scheme using force actuators and electret microphone sensors. There is a potential for dramatic cost savings, especially for large DMs beyond 1 m in diameter with thousands of actuators. Further, the suction cup attachment concept gives easy access for replacement of the actuators and microphones, which is important for maintenance.

The predicted DM performance using the SISO rate feedback and local position feedback approach described herein is promising and meets the requirement of an AO-system. The half power bandwidth of the transfer function between commanded actuator force and displacement of the faceplate for the simulation parameters used here is ~110 Hz, with less than 25° phase lag at a plausible AO bandwidth of 50 Hz. The performance-robustness of the proposed local control approach to uncertainty in the mirror structural dynamics is good, and better than either a purely collocated approach or a simple global model-based approach.

We have above assumed force actuators and position sensors to be ideal. Actual force actuators perform as low-pass filters whereas the position sensors based upon electret microphones can be seen as band-pass filters. The effect of high-frequency actuator and sensor dynamics will be small provided that the cut-off frequencies are beyond ~1.5 kHz where roll-off is intentionally introduced into the rate feedback in our model. The effect of actuator and sensor dynamics, together with implementation delays,

influences the phase margin for the rate feedback loop and thus will reduce the maximum tolerable separation between actuators and sensors in a quantifiable way. The influence of low-frequency microphone roll-off is also expected to be negligible, because the wavefront sensor for the outer AO loop provides sufficient information at low frequencies, where the microphones have low response.

Future work involves optimization of actuator and sensor design, new prototyping, including the modelled and measured dynamics of the force actuators and microphones in the model, and constructing a prototype mirror in the 1 m class.

The control architecture presented in this paper can also be used in other fields. For example, use of a very thin primary mirror in space telescopes is appealing for the weight budget; the controller developed here could be used to maintain the shape of the mirror.

References

1. Andersen T, Garpinger O, Owner-Petersen M, Bjoorn F, Svahn R, Ardeberg A. Novel concept for large deformable mirrors. *Opt Eng* 2006; 45(7): 073001113.
2. Benassi L, Elliott SJ. The equivalent impedance of power-minimizing vibration controllers on plates. *J Sound Vib* 2005; 283: 4
3. Biasi R, Andrighettoni M, Veronese D, et al. LBT adaptive secondary electronics. *Adap Opt Syst Technol II* 2003; 4839: 772782.
4. D'Andrea R, Dullerud G. Distributed control design for spatially interconnected systems. *IEEE Trans Autom Control* 2003; 48(9): 14781495.
5. Kendrew S, Doel P. Finite element analysis of carbon fibre composite adaptive mirrors. *Aviat Adap Opt* 2004; 5490: 15911599.
6. Kulkarni R, D'Andrea R, Brandl B. Application of distribution control techniques to the adaptive secondary mirror of Conrell's Large Atacama Telescope. *Proc SPIE* 2003; 4839:2003.
7. Lessard L, West M, MacMynowski D, Lall S. Warm-started wavefront reconstruction for adaptive optics. *J Opt Soc Am A* 2008; 25(5): 11471155.
8. Lyon RH. *Machinery Noise and Diagnostics*. Butterworth Publishing 1987.
9. MacMartin DG, Hall SR. Control of uncertain structures using an H₂ norm approach. *AIAA J Guid Control Dyn* 1991; 14(3):521–530.
10. MacMartin DG. Local, hierarchic, and iterative reconstructors for adaptive optics. *J Opt Soc Am* 2003; 20(6):10841093.
11. Miller DW, Grocott SCO. Robust control of the multiple mirror telescope adaptive secondary mirror. *Opt Eng* 1999; 38(8): 12761287.
12. Preumont A. *Vibration control of active structures: an introduction*. Academic Publishers, 2002.
13. Riccardi A, Brusa G, Salinari P, et al. Adaptive secondary mirrors for the large binocular telescope. *Astron Adap Opt Syst Appl* 2003; 5169: 159168.