

Slope estimation and wavefront reconstruction for laser tomography adaptive optics at the Giant Magellan Telescope

Marcos A. van Dam^a and Hajime Ogane^b

^aFlat Wavefronts, 21 Lascelles Street, Christchurch 8022, New Zealand

^bAdvanced Instrumentation and Technology Centre,
Research School of Astronomy and Astrophysics, the Australian National
University,
Mount Stromlo Observatory, Weston ACT 2611, Australia

ABSTRACT

The Laser Tomography Adaptive Optics (LTAO) system for the Giant Magellan Telescope employs six off-axis laser guide stars (LGSs) to estimate the on-axis wavefront. This paper addresses the challenges associated with wavefront sensing using highly elongated LGSs. We evaluate and compare several wavefront slope estimation algorithms currently baselined for the three Extremely Large Telescopes, highlighting their respective strengths and limitations. In particular, we consider their performance off-null operation, which occurs due both to laser jitter and to the use of centroid origins to compensate for non-common path aberrations. We show that both the weighted centroid and the matched filter algorithms perform well, with the former being more robust against LGS profile variations. Dynamic recentering of the spots further improves the centroiding accuracy. In the second part of the paper, we focus on tomographic wavefront reconstruction. We demonstrate that the applying accurate diagonal values to the noise covariance matrix reduces the wavefront error. This error can be further reduced by using a non-diagonal covariance matrix that incorporates the correlations between measurements within a subaperture. End-to-end Monte-Carlo simulations are run using CEO, an optical propagation package developed for the GMT.

Keywords: adaptive optics, laser tomography adaptive optics, centroiding, wavefront sensing, wavefront reconstruction, Giant Magellan Telescope

1. INTRODUCTION

The three Extremely Large Telescopes (ELTs) currently under design or construction will all employ Laser Tomography Adaptive Optics (LTAO) systems to provide high Strehl ratio correction over a small field of view. LTAO systems use multiple off-axis laser guide stars (LGSs) to estimate the on-axis wavefront. The LGSs are propagated from either behind the secondary mirror in the case of the Thirty Meter Telescope (TMT) or from the side of the telescope in the case of the Giant Magellan Telescope (GMT) and of the Extremely Large Telescope (ELT).

The LGS spots are elongated due to the finite thickness of the sodium layer in the mesosphere, which is located at an altitude of approximately 80km to 100km above sea level. All planned LTAO systems for the ELTs will measure the wavefront using Shack-Hartmann wavefront sensors (SH WFSs). The SH WFS subdivides the pupil using a lenslet array, and each subaperture produces an image of the LGS. The effect of LGS elongations on the WFS measurements has been studied extensively for the two adaptive optics (AO) systems at the W. M. Keck Observatory.[1, 2] Numerical modeling of the elongations has also been implemented for the TMT, and the methods described in Wang *et al.* are adopted in this paper.[3]

The LGS elongation is proportional to the distance in the pupil plane between the launch telescope and the subaperture. Due to their large primary mirror diameters, the ELTs will have highly elongated LGS spots for subapertures situated far from the launch telescope, and this has long been considered a showstopper.[4] The

Send correspondence to Marcos van Dam, e-mail: marcos@flatwavefronts.com

advent of fast, large format, low read noise detectors,[5] as well as lasers that produce a high sodium return,[6] has enabled the design of LTAO systems for the ELTs.

The use of this technology alone is not sufficient to reduce the measurement noise to acceptable levels. The displacement of the spots on the WFS is proportional to the wavefront slope. The slopes are conventionally estimated using a centroid (center of gravity) algorithm, but due to the large number of pixels required to sample the elongated spots adequately, the measurement noise using the centroid is unacceptably high.[7] In this paper, we describe the wavefront slope estimation algorithms planned for the ELTs and evaluate their performance and their dynamic range. In addition, we show that incorporating full (non-diagonal) noise covariance matrices in the wavefront reconstructors can reduce the impact of the increased noise along the elongation direction while making full use of the measurements in the orthogonal direction.

The remainder of the paper is as follows. In Section 2, we describe the wavefront slope estimation algorithms proposed by the three ELTs. Section 3 describes the end-to-end simulations used to evaluate the performance of the centroiding algorithms using simulation parameters representative of the GMT LTAO system. The simulation results are presented in Section 4, comparing the various algorithms with respect to computation time, linearity, sensitivity to measurement noise and sensitivity to uplink tip-tilt errors. Section 5 presents the tomographic wavefront reconstruction algorithms. In particular, we show that a non-diagonal noise covariance matrix leads to a reduced wavefront error. Finally, conclusions are drawn in Section 6.

2. WAVEFRONT SLOPE ESTIMATION

In this section, we review the wavefront slope estimation algorithms proposed for the ELTs. Before applying a centroiding algorithm, all subaperture images are background subtracted and thresholded to zero to remove any negative pixel values.

2.1 Centroid

First, we present the centroid, also known as the center of gravity (CoG) algorithm, which is the “gold standard” for wavefront slope estimation. The centroid, (x_c, y_c) is computed as:

$$\begin{aligned} x_c &= \frac{\sum_x x \sum_y I(x, y)}{\sum_x \sum_y I(x, y)} \\ y_c &= \frac{\sum_y y \sum_x I(x, y)}{\sum_x \sum_y I(x, y)} \end{aligned} \quad (1)$$

where $I(x, y)$ is the intensity of the pixel at position (x, y) .

While the centroid algorithm suffers from high noise when there is a large number of pixels per subaperture,[7] as is the case for elongated LGS spots, it is useful as a comparison for noiseless wavefront sensor measurements.

2.2 Weighted centroid

The weighted centroid (or weighted center of gravity, WCoG) algorithm is a modification of the standard centroid algorithm that applies a weighting function to the pixel intensities. The weighting function is typically chosen to match the expected shape of the LGS spot, which can help reduce the impact of noise, especially when there are a large number of pixels. The weighted centroid is computed in the following manner:


$$\begin{aligned} x_c &= \frac{\sum_x x \sum_y w(x, y) I(x, y)}{\sum_x \sum_y w(x, y) I(x, y)} \\ y_c &= \frac{\sum_y y \sum_x w(x, y) I(x, y)}{\sum_x \sum_y w(x, y) I(x, y)} \end{aligned} \quad (2)$$

where $w(x, y)$ represents the pixel weights applied to the pixel intensities. The pixel weights vary from subaperture to subaperture, depending on the elongation of the LGS spot. They can also vary with time as the elevation or azimuth of the telescope change or the sodium layer profile changes.

The weighted centroid algorithm is baselined for the HARMONI instrument at the ELT,[8] and is also favored for the LTAO system at the GMT. The choice of weighting function has a strong effect on the performance. The weighting functions can be computed using a model of the sodium layer profile, or they can be derived from the measured LGS spot images themselves. The latter approach requires averaging multiple frames to reduce noise, which can limit the algorithm’s ability to track rapid changes in the sodium layer profile.

The modeling for HARMONI assumes that the LGS elongation is well known. The weighting function is computed by convolving the LGS elongation by a Gaussian kernel representing the LGS spot. The LGS spot used to compute the weighting function is much larger (typically a factor of 2.5) than the real LGS, as shown in Figure 1(b). The weights are thus a smoother version of the expected intensity at each subaperture.

The GMT has chosen a different approach for the pixel weights. Initially, a binary weights map was used, where the value of the weight is unity for pixels with an expected intensity greater than half the read noise, and zero elsewhere. This works very well for the case where the spots are well centered on the subaperture, but the performance degrades if the spots are off-center. We have added the option of a soft weights map, where the weights are defined as the expected intensity divided by a quarter of the read noise, with a maximum value of unity. This approach retains more information when the spots are off-center, as shown in Figure 1(d).



(a) Noiseless subaperture (b) HARMONI weights (c) Hard weights (d) Soft weights
Figure 1: Typical noiseless elongated spot and corresponding weights used in the weighted centroid.

2.3 Weighted centroid with dynamic recentering

The weighted centroid with dynamic recentering (WCoG-DR) algorithm is a modification of the weighted centroid algorithm that includes an iterative recentering step. This approach is designed to improve the accuracy of the slope estimates when the LGS spots are not well centered on the subapertures, which occurs due to uplink tip-tilt errors or large centroid origins. The algorithm is implemented as follows:

1. Compute the weighted centroid of the subimage.
2. Displace the subimage based on the computed weighted centroid.
3. Compute the weighted centroid of the displaced subimage and add to the weighted centroid of Step 1.

In this work, the subimage is only recentered once, but the process can be iterated multiple times if necessary. The subimage is displaced using a Fourier shift:

1. Compute the Fourier transform of the subimage using the FFT algorithm.
2. Multiply the Fourier transform by a phase ramp corresponding to the desired subimage shift.
3. Compute the inverse Fourier transform to obtain the shifted subimage.

2.4 Matched filter

The matched filter was introduced as a practical solution to the spot elongation problem facing TMT by Gilles and Ellerbroek.[9] The principle behind the matched filter is to compute a derivative of the normalized pixel intensities with respect to displacements in the x and y directions. It produces low-noise slope estimates for elongated LGS spots. However, the linear range of the matched filter is very limited and low-order aberrations result when the spots are decentered by even a fraction of a pixel. An elegant solution to this problem is to add a set of linear constraints corresponding to a displacement by a pixel in each direction.[10]

For each $N_{\text{pix}} \times N_{\text{pix}}$ pixel subaperture image, the matched filter computes the slope as follows by multiplying each pixel in the image by a $2 \times N_{\text{pix}}^2$ matrix. As a result, computing the centroid using the matched filter is computationally very efficient. Details on the implementation of the matched filter and how to compute the matrix are found in Gilles and Ellerbroek.[10] The matched filter is baselined for the NFIRAOS system at TMT.[3]

3. SIMULATION DESCRIPTION

In this section, we describe the simulations performed to evaluate the performance of the proposed centroiding methods. We used end-to-end simulations in the CEO environment,[11] a simulation tool written specifically to model the optical behavior of the GMT. The simulations were written in Python and use GPU processing for speed with the aid of the CuPy library.

3.1 Disturbances

The only disturbances in these simulations are atmospheric turbulence. The atmospheric parameters are derived from the typical-typical profile for January 2008 from Goodwin[12] and are reproduced in Table 1. This turbulence

Table 1: Turbulence profile used in the simulations.

Elevation (m)	Turbulence fraction	Wind speed (m/s)	Wind direction (°)
25	0.126	5.65	0.78
275	0.087	5.80	8.25
425	0.067	5.89	12.48
1250	0.350	6.64	32.50
4000	0.227	13.29	72.10
8000	0.068	34.83	93.20
13000	0.075	29.42	100.05

profile has a higher fraction of turbulence in the free atmosphere compared to profiles from other studies, including those measured by the same author four months earlier. This leads to strong angular anisoplanatism, which is compensated in the tip-tilt sensor using a dedicated DM. Unless otherwise specified, the value of r_0 is 0.16 m at a wavelength of 500 nm at zenith with an outer scale of 25 m, median values for the site. All of the simulations were run at zenith, because that is where the elongation is at its greatest.

No telescope windshake or other telescope aberrations, including segment piston errors, are included in these simulations, in order to accentuate the difference between the different sensing algorithms.

3.2 Sodium profiles

The sodium profiles used in this report originate from the University of British Columbia (UBC) Lidar and reported in multiple papers by Pfrommer and Hickson.[13, 14] There are 176 data sets where there are 180 measurements of the intensity profiles every 10s with a vertical resolution of 200 m. During the course of our work, all of the sodium profiles were evaluated and sorted by the effect of the LGS aberrations on the wavefront error. Unless otherwise specified, the sodium profile used in these simulations is the typical profile shown in Figure 2.




Figure 2: Sodium profile used in the simulations.

The temporal evolution of this sodium profile is shown in Figure 3.




Figure 3: Sodium profile evolution as a function of time (in seconds).

The sodium density data reported by the University of British Columbia LIDAR are intensity profiles of the fluorescent sodium layer viewed from the ground and are not true sodium density profiles. This makes the computation of the LGS elongation easier, because we do not need to take into account the distance to the sodium layer when computing the intensity.

3.3 Wavefront sensors

Two wavefront sensing units are simulated: the laser tomography WFS (LTWS) and the tip-tilt sensor.

The LTWS consists of six Shack-Hartmann WFS with 60×60 lenslets across the pupil. The LGSs are projected in a regular hexagonal constellation on the sky at an off-axis distance of $30''$. Physical optics are used to model the spots on both of the WFSs. The relevant noise parameters for the LGS WFSs are tabulated in Table 2.

Table 2: LGS WFS parameters used in the simulations.

Parameter	Value
Wavelength	589 nm
Bandwidth	1 nm
Optical throughput	48%
Laser power	22 W
Number of LGSs	6
Sodium column density	$4 \times 10^{13} \text{ m}^{-2}$
Intrinsic LGS spot size	1'' FWHM
Subapertures across pupil	60×60
Pixels per subaperture	14×14
Pixel scale	0.857''
Frame rate	500 Hz
Readout noise	$2.5 e^-$
Dark current	$0 e^-/s$
Sky background	$0 e^-/s$
Readout bits	12
ADU gain	15 ADU/ e^-

The detector for the tip-tilt sensor is assumed to be the SAPHIRA, an infrared e-APD array manufactured by Leonardo. [15, 16] Table 3 tabulates the photometry and noise parameters for the tip-tilt sensor. An on-axis guide star with a K-band magnitude of 14 is used.

Table 3: Photometric and noise parameters for the tip-tilt sensor used in the simulations.

Parameter	Value
Central wavelength	2.179 μm
Bandwidth	0.41 μm
Photometric zero point	7.0×10^{11}
Sky background (magnitude / arcsec ²)	12.71
Optical throughput	60%
Quantum efficiency	72%
Excess noise factor	1.35
Read noise	$0.4 e^-$
Dark current	$10 e^-$

3.4 Modeling the LGS elongation

In this section, we describe how the LGS elongation is modeled in the simulations. The LGS appears elongated due to the parallax effect. For simplicity, let us first consider an observation at zenith with the sodium density centered around height h_0 . We propagate an LGS from a laser launch telescope at an altitude of the of 2282 m above sea level, and image the return using a particular subaperture separated from the launch telescope by a distance b . The three locations for the launch telescopes are shown in Figure 4. The value of b is 1.83 m for the least elongated spots, and 23.60 m for the most elongated spots.




Figure 4: Location of the laser launch telescopes.

The WFS is nominally conjugate to h_0 . Ignoring any atmospheric or telescope aberrations, a ray of light stemming from a sodium atom at height h_0 will be imaged at the center of the subaperture. For a sodium atom at height h_i , the angular offset is $\theta_i - \theta_0$ (Figure 5).




Figure 5: Relationship between the height of a sodium atom and the angle it subtends on the wavefront sensor.

The procedure for computing the LGS elongation transfer function is as follows. First, the sodium density as a function of altitude is loaded from a file. Next, a grid of pixels is created with the resolution of the SH WFS in CEO. In this case, we are using 60x60 subapertures across the pupil, 8 pixels per subaperture in the pupil and an oversampling parameter of 2 at a wavelength of 589 nm, so the unbinned pixel scale is $0.1429''$. All of the calculations are performed at this resolution. We select an array of 84×84 pixels, which will be binned at the final step to produce 14×14 binned pixels with a binned pixel scale of $0.8576''$. For each subaperture, we compute the distance and angle to the laser launch telescope.

If we write the components of b as b_x and b_y , then for each layer i with height h_i we compute the displacements x_i and y_i using the relationships

$$x_i = \tan(h_i/b_x) - \tan(h_0/b_x) \quad (3)$$

and

$$y_i = \tan(h_i/b_y) - \tan(h_0/b_y) \quad (4)$$

The PSF corresponding to that layer is computed by applying a subpixel shift using bilinear interpolation to a delta function. This PSF is multiplied by the relative intensity of the corresponding sodium altitude and coadded to all of the other sodium altitudes.

3.5 Wavefront control

The simulations were run for a total of 1.5s at a rate of 500 Hz. The feedback control loops were modeled as an integral controller with a loop gain of 0.6. There was also a pure delay of one cycle, which was added to the one cycle inherent in the camera stare and the zero-order hold of the wavefront corrector. The high-order reconstructor consists of a minimum-variance tomographic reconstructor operating in pseudo open-loop. The tip-tilt loop is controlled using a least-squares reconstructor in closed-loop. The adaptive secondary mirror (ASM) is in the common path and corrects the wavefront incident on the science instrument as well as all of the WFSs. The first 200 Karhunen-Loève modes are corrected for each of the seven segments, for a total of 1400 modes. The temporal dynamics of the ASM and were deemed negligible for this study.

4. SIMULATION RESULTS

In this section, we run simulations to compare the different centroiding algorithms. The following properties are compared: computation time, sensitivity to measurement noise, sensitivity to uplink tip-tilt errors and sensitivity to LGS profile variations.

4.1 Computation time of centroiding algorithms

An important concern for the ELTs is whether the centroiding algorithms can be implemented in real time. The algorithms were implemented in Python using the CuPy library to take advantage of GPU processing. The timing results were obtained using an NVIDIA RTX A5000 GPU and are tabulated in Table 4. These numbers are sufficiently low to assume that using a more powerful GPU and a more efficient implementation would allow real-time operation at 500 Hz to 1000 Hz.

Table 4: Timing results for the different centroiding algorithms.

Algorithm	Total time (ms)
Centroid	1.1
Weighted centroid	1.2
Weighted centroid with dynamic recentering	5.4
Matched filter	1.3

4.2 Linearity

The linearity of the centroiding algorithms was evaluated by applying a tip of $0.1''$ to each WFS and recording the change in the centroids. For each of the six WFSs, a different algorithm was used, and the response is plotted in Figure 6. The results show that the centroid and matched filter are extremely linear. The soft window is significantly more linear than the hard window, but dynamically recentering the spots alleviates the non-linearity of the hard window. The HARMONI weights produce a response that depends on the elongation. We could calibrate the centroid gain on a subaperture-by-subaperture basis, but this would be cumbersome and error-prone in real life. Instead, the centroids were multiplied by a constant gain of 1.5 so that the loop gain was approximately equal for all centroiding algorithms.




Figure 6: Response of each of the centroiding algorithms to a tip of $0.1''$.

4.3 Measurement noise

The effect of measurement noise has two components: a bias in the centroids, and a random, approximately Gaussian component. The measurement noise is propagated via the wavefront reconstructor, which leads to a residual wavefront error. The impact of measurement errors is not uniform, since in LTAO some measurements are more redundant than others.

Simulations were run with and without measurement noise, with the results displayed in Table 5. When there is no noise, all of the algorithms produce similar results except for the weighted centroid with the HARMONI weights. In the presence of noise, the conventional centroid ceases to be competitive. It is encouraging that there are several feasible algorithms. The soft window is better than the hard window, and dynamic recentering improves the performance of both.

Table 5: Wavefront error (nm) for the different centroiding algorithms with and without noise.

Algorithm	No noise	With noise
Centroid	180.0	249.8
Weighted centroid, hard window	186.7	200.1
Weighted centroid, soft window	180.2	196.9
Weighted centroid, HARMONI weights	212.3	224.0
Weighted centroid, hard window, dynamic recenter	182.2	197.2
Weighted centroid, soft window, dynamic recenter	178.7	193.8
Matched filter	178.8	200.0

4.4 Uplink tip-tilt errors

In Section 4.3, the spots were well centered around their nominal positions. However, in practice, the spots will be displaced due to uplink tip-tilt errors. Measured residual uplink tip-tilt time-series from the adaptive optics facility (AOF) at the VLT were scaled and applied as a tip-tilt to the wavefront experienced by the LGS WFSs. The GMT is required to deliver a residual uplink tip-tilt of less than $0.15''$ RMS per axis. The results for various levels of uplink tip-tilt errors are tabulated in Table 6. Measurement noise was applied to all of the simulations.

Table 6: Wavefront error (nm) for the different centroiding algorithms with various levels of uplink tip-tilt errors.

Algorithm	0''	0.05''	0.1''	0.15''	0.2''
Weighted centroid, hard window	200.1	200.6	205.7	213.1	222.4
Weighted centroid, soft window	196.9	198.0	200.8	205.5	211.2
Weighted centroid, HARMONI weights	224.0	228.5	239.0	256.2	276.5
Weighted centroid, hard window, dynamic recenter	197.2	197.0	197.1	197.5	197.8
Weighted centroid, soft window, dynamic recenter	193.8	193.9	195.6	195.4	195.3
Matched filter	200.0	200.9	203.2	206.8	211.5

4.5 LGS profile variations

The simulations to date assume that we know the sodium profile perfectly. In practice, the sodium profile varies with time, and at best we have an estimate of the profile at some time in the past. Simulations were run where the weights were computed using a profile that is 100 s old. In addition, the simulations were repeated with quite a different turbulence profile, as shown in Figure 7. The simulations were all run with measurement noise and 0.15'' of uplink tip-tilt errors. The results, tabulated in Table 7, suggest that the weighted centroid is largely insensitive to changes in the sodium profile. However, the matched filter requires an excellent estimate of the sodium profile to perform well. The matched filter uses the derivative with respect to x - and y of the expected spot intensity to estimate the wavefront slope, and changes in sodium structure lead to changes in the derivative. By contrast, the weighted centroid only needs to know the expected location of the spot, and not its expected intensity distribution.




Figure 7: Sodium profile used to compute the weights for the “incorrect” case.

Table 7: Wavefront error (nm) with different assumed sodiums profile for the different centroiding algorithms.

Algorithm	Correct	100 s stale	Incorrect
Weighted centroid, hard window	213.1	213.0	212.4
Weighted centroid, soft window	205.5	205.4	206.5
Weighted centroid, HARMONI weights	256.4	251.0	238.3
Weighted centroid, hard window, dynamic recenter	197.5	196.9	198.4
Weighted centroid, soft window, dynamic recenter	195.4	196.8	197.2
Matched filter	206.8	207.2	231.3

5. WAVEFRONT RECONSTRUCTION

5.1 Minimum-variance tomographic reconstruction

The wavefront reconstruction is performed using a minimum-variance reconstructor operating in pseudo open-loop. The reconstructor uses a model of the atmospheric turbulence, telescope aberrations and measurement

noise to compute the optimal tomographic wavefront estimate. The tomographic reconstructor computes the wavefront along a grid of points in the direction of the science target, and then projects the wavefront onto the modes of the adaptive secondary mirror (ASM). The reconstructor, R , is computed as:

$$R = P_{asm} C_{xs} (C_{ss} + C_{nn})^{-1} \quad (5)$$

where P_{asm} is the projection matrix onto the ASM modes, C_{xs} is the cross-covariance matrix between the wavefront at the science target and the WFS measurements, C_{ss} is the covariance matrix of the WFS measurements and C_{nn} is the noise covariance matrix of the WFS measurements. The matrices C_{xs} and C_{ss} are computed using a model of the atmospheric turbulence in open loop. The covariance matrices are computed using code developed by Vidal *et al.*[17]

Since the LTAO system works in closed loop, the open-loop slopes need to be computed from the closed-loop slopes and the ASM commands. Denoting the matrix that converts ASM commands to slopes as M_{asm} , the pseudo open-loop slopes, s_{pol} are computed as:

$$s_{pol} = s + M_{asm} a_{asm} \quad (6)$$

where a_{ASM} are the ASM commands, and the residual wavefront estimate, u_{asm} , is:

$$\begin{aligned} u_{asm} &= R s_{pol} - a_{asm} \\ &= R s + (R M_{asm} - I) a_{asm} \\ &= R s + D a_{asm} \end{aligned} \quad (7)$$

where $D = R M_{asm} - I$ and I is the identity matrix.

5.2 Theory of the non-diagonal noise covariance matrix

To our knowledge, every existing AO system uses a diagonal noise covariance matrix, where the noise from each measurement is assumed to be uncorrelated. Some measurements have a higher covariance, for example due to partially illuminated subapertures or elongated LGS spots. However, a diagonal covariance matrix does not incorporate correlations between different measurements. In the case of elongated LGS spots, the noise in the direction of elongation is higher than the noise in the perpendicular direction, even if these axes do not align with the x - and y -axes of the WFS. As a result, the noise in different measurements is correlated.

Let us consider the following toy problem, with a single subaperture used to measure the location of a single LGS spot, which can be accurately modeled as a Gaussian. Four different cases are considered: a 1.5'' FWHM unelongated spot, a 3'' FWHM spot elongated along the x -axis, and the same spot rotated by 45° and 30°, as shown in Figure 8.




Figure 8: LGS spots used in the toy problem with a single subaperture.

In each case, we compute the covariance of the measurement noise in arbitrary but consistent units (Table 8).

Table 8: Noise covariance matrices for elongated spots.

$$\begin{array}{cccc} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} & \begin{bmatrix} 4 & 0 \\ 0 & 1 \end{bmatrix} & \begin{bmatrix} 2.50 & -1.50 \\ -1.50 & 2.50 \end{bmatrix} & \begin{bmatrix} 3.20 & -1.27 \\ -1.27 & 1.74 \end{bmatrix} \\ \text{(a)} & \text{(b)} & \text{(c)} & \text{(d)} \end{array}$$

5.3 Simulations using a non-diagonal noise covariance matrix

The noise covariance matrix was computed experimentally in simulations by making 1000 measurements of the centroids in the presence of measurement noise, and then computing the covariance of the measurements. Three covariance matrices were computed:

- A diagonal covariance matrix with uniform entries consistent with the typical noise variance.
- A diagonal covariance matrix where the values represent the variance of each wavefront slope.
- A non-diagonal covariance matrix where the covariance of both measurements from the same subaperture are included and all values that pertain to measurements from unrelated subapertures are set to zero.

The reconstructor was recomputed using the three noise covariance matrices and the simulations were run, with the results tabulated in Table 9. A weighted centroid with a soft window but no dynamic recentering was used. The results show that there is a significant benefit to using a covariance matrix with the right values along the leading diagonal. In addition, a measurable improvement is obtained by using a non-diagonal noise covariance matrix. This is especially true for the case with noise and no jitter, which is the case for which the covariance matrix was computed. Even when there is no measurement noise, a reconstructor with information about the elongation direction performs better.

Table 9: Wavefront error (nm) for the different using a diagonal and non-diagonal noise covariance matrix. A weighted centroid with the soft window but no dynamic recentering was used.

	Uniform	Diagonal	Non-diagonal
No noise, no jitter	180.2	178.6	177.6
With noise, no jitter	196.9	187.2	184.3
With noise, with 0.15'' jitter	205.5	196.7	195.3

6. CONCLUSION

In this paper, we discuss a number of different centroiding algorithms for elongated LGS spots. The conventional center-of-mass algorithm exhibits an excellent linear response and excellent performance in the absence of noise, but the performance degrades significantly in the presence of measurement noise. The noise can be mitigated by using a weighted centroid with a soft window, where the weight values are unity where high signal-to-noise is expected, tapering to zero where the signal-to-noise is low. The performance can be further improved by dynamically recentering the spots before recomputing the weighted centroid. This is beneficial because the centroid weights bias the centroids towards the center of the subaperture. With these two modifications to the centroid algorithm, the performance is excellent even in the presence of uplink tip-tilt errors and changes in the sodium profile. An alternative approach, selected by the TMT, is to use a matched filter. The matched filter is very linear, has a large dynamic range and exhibits excellent performance in the presence of measurement noise, but it is very sensitive to changes in the sodium profile. Both the matched filter and the weighted centroid are very efficient and suitable for real-time operation. The additional step of dynamically recentering the spots adds some computational complexity, but it is still feasible to implement in real time.

We find that for an LGS WFS with elongated spots, using the use of an noise covariance matrix that incorporates the variance of each measurement leads to a significant reduction in wavefront error. The error can be further reduced by using a non-diagonal noise covariance matrix that includes the correlations between the x - and y -measurements from each subaperture. This is a new result, and to our knowledge, no existing AO system uses a non-diagonal noise covariance matrix.

REFERENCES

- [1] Van Dam, M. A., Bouchez, A. H., Le Mignant, D., and Wizinowich, P. L., “Quasi-static aberrations induced by laser guide stars in adaptive optics,” *Optics Express* **14**(17), 7535–7540 (2006).
- [2] Clare, R. M., van Dam, M. A., and Bouchez, A. H., “Modeling low order aberrations in laser guide star adaptive optics systems,” *Optics Express* **15**(8), 4711–4725 (2007).
- [3] Wang, L., Boyer, C., Ellerbroek, B., Herriot, G., and Veran, J.-P., “Simultaneous sodium profile estimation and LGS SH-WFS pixel processing optimization using LGS sub-aperture images,” *Journal of Astronomical Telescopes, Instruments, and Systems* **9**(1), 019002–019002 (2023).
- [4] Gavel, D., “Technology challenges for adaptive optics on extremely large telescopes,” in [*Beyond conventional adaptive optics: a conference devoted to the development of adaptive optics for extremely large telescopes. Proceedings of the Topical Meeting held May 7-10, 2001, Venice, Italy. Edited by E. Vernet, R. Ragazzoni, S. Esposito, and N. Hubin. Garching, Germany: European Southern Observatory, 2002 ESO Conference and Workshop Proceedings, Vol. 58, ISBN 3923524617, p. 47*], **58**, 47 (2002).
- [5] Gach, J.-L., Boutolleau, D., Brun, C., Carmignani, T., Clop, F., Feautrier, P., Florentin, M., Pettigiani, J., Stadler, E., Tugnoli, J., et al., “C-BLUE one: a new CMOS camera dedicated for laser guide star wavefront sensing on ELTs,” in [*Adaptive Optics Systems VII*], **11448**, 114483C, SPIE (2020).
- [6] Enderlein, M., “SodiumStar—the sodium guide star laser by TOPTICA Projects GmbH,” in [*SPIE Exhibition Product Demonstrations*], AS20EX0F, SPIE (2020).
- [7] Thomas, S., Fusco, T., Tokovinin, A., Nicolle, M., Michau, V., and Rousset, G., “Comparison of centroid computation algorithms in a Shack–Hartmann sensor,” *Monthly Notices of the Royal Astronomical Society* **371**(1), 323–336 (2006).
- [8] Heritier, C. T., Neichel, B., Fusco, T., Sauvage, J.-F., Correia, C., Thatte, N., Madec, P.-Y., Paufigue, J., Le Louarn, M., Morris, T., et al., “HARMONI at ELT: status of the LTAO mode,” in [*Adaptive Optics Systems IX*], **13097**, 1309721, SPIE (2024).
- [9] Gilles, L. and Ellerbroek, B., “Shack–Hartmann wavefront sensing with elongated sodium laser beacons: centroiding versus matched filtering,” *Applied optics* **45**(25), 6568–6576 (2006).
- [10] Gilles, L. and Ellerbroek, B., “Constrained matched filtering for extended dynamic range and improved noise rejection for Shack–Hartmann wavefront sensing,” *Optics letters* **33**(10), 1159–1161 (2008).
- [11] Conan, R., “Cuda Engined Optics.” <https://github.com/rconan/CEO> (2025).
- [12] Goodwin, M. S. et al., “Turbulence profiling at siding spring and las campanas observatories,” (2009).
- [13] Pfrommer, T. and Hickson, P., “High-resolution lidar observations of mesospheric sodium and implications for adaptive optics,” *Journal of the Optical Society of America A* **27**(11), A97–A105 (2010).
- [14] Pfrommer, T. and Hickson, P., “High resolution mesospheric sodium properties for adaptive optics applications,” *Astronomy & Astrophysics* **565**, A102 (2014).
- [15] Atkinson, D., Hall, D., Baranec, C., Baker, I., Jacobson, S., and Riddle, R., “Observatory deployment and characterization of SAPHIRA HgCdTe APD arrays,” in [*High Energy, Optical, and Infrared Detectors for Astronomy VI*], **9154**, 405–416, SPIE (2014).
- [16] Finger, G., Baker, I., Alvarez, D., Ives, D., Mehrgan, L., Meyer, M., Stegmeier, J., and Weller, H. J., “SAPHIRA detector for infrared wavefront sensing,” in [*Adaptive Optics Systems IV*], **9148**, 427–442, SPIE (2014).
- [17] Vidal, F., Gendron, E., and Rousset, G., “Tomography approach for multi-object adaptive optics,” *Journal of the Optical Society of America A* **27**(11), A253–A264 (2010).