

Fast & Furious focal-plane wavefront sensing

Visa Korkiakoski,^{1,2,*} Christoph U. Keller,² Niek Doelman,^{2,3} Matthew Kenworthy,²
Gilles Otten,² and Michel Verhaegen¹

¹Delft Center for Systems and Control, Mekelweg 2, 2628CD Delft, The Netherlands

²Leiden Observatory, Leiden University, P.O. Box 9513, 2300 RA Leiden, The Netherlands

³TNO Technical Sciences, Stieltjesweg 1, 2628CK Delft, The Netherlands

*Corresponding author: v.a.korkiakoski@tudelft.nl

Received 7 March 2014; revised 26 May 2014; accepted 2 June 2014;
posted 3 June 2014 (Doc. ID 207889); published 9 July 2014

We present two complementary algorithms suitable for using focal-plane measurements to control a wavefront corrector with an extremely high-spatial resolution. The algorithms use linear approximations to iteratively minimize the aberrations seen by the focal-plane camera. The first algorithm, Fast & Furious (FF), uses a weak-aberration assumption and pupil symmetries to achieve fast wavefront reconstruction. The second algorithm, an extension to FF, can deal with an arbitrary pupil shape; it uses a Gerchberg–Saxton (GS)-style error reduction to determine the pupil amplitudes. Simulations and experimental results are shown for a spatial-light modulator controlling the wavefront with a resolution of 170×170 pixels. The algorithms increase the Strehl ratio from ~ 0.75 to 0.98 – 0.99 , and the intensity of the scattered light is reduced throughout the whole recorded image of 320×320 pixels. The remaining wavefront rms error is estimated to be ~ 0.15 rad with FF and ~ 0.10 rad with FF-GS. © 2014 Optical Society of America

OCIS codes: (010.1080) Active or adaptive optics; (120.5050) Phase measurement; (280.4788) Optical sensing and sensors; (220.1080) Active or adaptive optics.

<http://dx.doi.org/10.1364/AO.53.004565>

1. Introduction

When an object is imaged, variations of the refractive index in the medium, as well as optical alignment and manufacturing errors, distort the recorded image. This problem is typically solved using active or adaptive optics, where a deformable mirror, spatial-light modulator (SLM), or a comparable device corrects the propagating wavefront (WF). Typically, such systems are built with a separate optical arm to measure the distorted WF because extracting the WF information from only focal-plane images is not trivial. However, focal-plane WF sensing is an active topic—not only to simplify the optical design but also to eliminate the noncommon path aberrations limiting the performance of high-contrast adaptive-optics systems.

The most popular method for focal-plane WF sensing is perhaps the Gerchberg–Saxton (GS) error reduction algorithm [1,2] and its variations, for instance [3,4]. These are numerically very efficient algorithms, easily modified for different applications. However, they suffer from lack of accuracy, in particular because their iterative improvement procedure often stagnates at a local minimum.

Various alternatives have been proposed, and a popular approach is to use general numerical optimization techniques to minimize an error function; examples include [5–7]. However, when the number of optimization parameters is increased, the computational requirements generally rise unacceptably fast. The high computational costs are problematic, for instance, in astronomy; the largest future adaptive-optics system is envisioned to have a WF corrector of a size of 200×200 elements [8].

The numerical issues can be significantly reduced if the unknown WF is sufficiently small. This is the case, for example, when calibrating the noncommon path aberrations. Previous works have exploited small-phase approximations [9–12], but the implementations are generally not easily extended to WF correction at extremely large resolution, such as over 100×100 elements.

In this paper, we present two algorithms capable of extremely fast control of a WF correcting device with 20,000–30,000 degrees of freedom.

The first algorithm, Fast & Furious (FF), has been published before [13–15]. It relies on small WF aberrations, pupil symmetries, and phase diversity to achieve very fast WF reconstruction. However, FF approximates the pupil amplitudes as an even function that does not necessarily exactly match the real situation.

To improve the WF correction beyond the accuracy of FF, a natural way is to use approaches similar to the GS algorithm. However, the standard modifications of the algorithm are sensitive to the used phase diversities, in particular when the pupil amplitudes are not known, and they do not work with iterative WF correction as in FF. Therefore, our second algorithm combines FF and GS in a way that can be used not only to correct the WF, but also to estimate the pupil amplitudes—for which we make no assumptions. This comes at a cost in terms of noise sensitivity and instabilities as well as more demanding computational requirements.

At first, we illustrate the motivation and principles of the FF algorithm in Section 2. Next, Section 3 describes the Fast & Furious Gerchberg–Saxton (FF-GS) algorithm in detail. Section 4 describes the used hardware; Section 5 shows simulation and experimental results; and Section 6 draws the conclusions.

2. Fast & Furious

The FF algorithm is based on iteratively applying a weak-phase approximation of the WF. The main principle of the weak-phase solution is presented in [16], but we found slight modifications [13], leading to significantly better performance. The algorithm uses focal-plane images and phase-diversity information to solve the WF, and the estimated WF is corrected with a WF correcting device. The

correction step produces phase-diversity information and a new image that are again used to compute the following phase update. The schematic illustration of the algorithm is shown in Fig. 1.

An important aspect of the algorithm is to maximize the use of the most recent point-spread function (PSF)—denoted as Image 1 in Fig. 1. In the weak-phase regime, a single image is sufficient to estimate both the full odd- WF component and the modulus of the even component of the focal-plane electric field. The phase diversity is needed only for the sign determination since we assume the WF aberrations are small. This makes the FF substantially less prone to noise and stability issues as compared to approaches relying more on the phase-diversity information—such as the FF-GS.

Section 2.A explains the details of the weak-phase solution, and Section 2.B discusses the practical aspects when implementing the algorithm.

A. Weak-Phase Solution

A monochromatic PSF can be described by Fraunhofer diffraction and is given by the squared modulus of the Fourier transform of the complex electric field in the pupil plane

$$p = |\mathcal{F}\{A \exp(i\phi)\}|^2, \quad (1)$$

where A is the pupil amplitude describing transmission and ϕ is the WF in the pupil plane.

The second-order approximation of the PSF, in terms of the WF expansion, can be written as

$$p = |\mathcal{F}\{A + iA\phi - 0.5A\phi^2\}|^2. \quad (2)$$

The phase ϕ can be represented as a sum of even and odd functions

$$\phi = \phi_e + \phi_o, \quad (3)$$

and Eq. (2) can then be written as

$$p = |\mathcal{F}\{A + iA\phi_e + iA\phi_o - 0.5A\phi_e^2 - 0.5A\phi_o^2 - A\phi_e\phi_o\}|^2. \quad (4)$$

We make the assumption that A is even, and therefore all the terms here are either even or odd. Thus, the corresponding Fourier transforms are then either

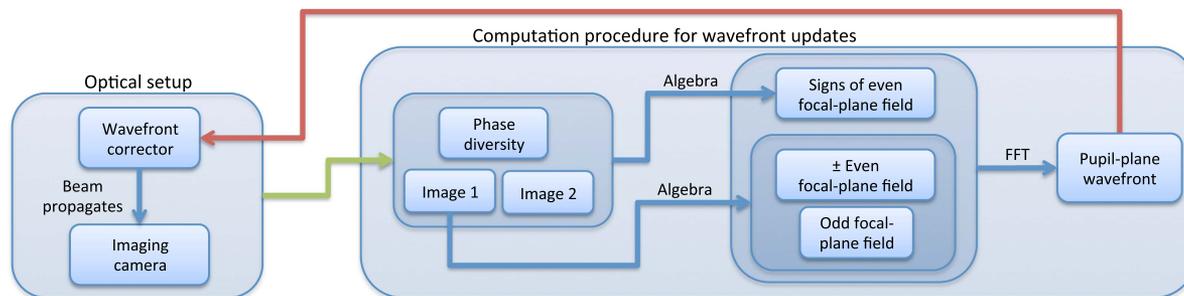


Fig. 1. Schematic illustration of the FF algorithm.

Table 1. Notations and Symmetries

Aperture Plane			Fourier Plane		
Term	Re/Im	Symmetry	Term	Re/Im	Symmetry
A	Real	Even	a	Real	Even
$A\phi_e$	Real	Even	v	Real	Even
$A\phi_o$	Real	Odd	iy	Imaginary	Odd
$A\phi_e^2$	Real	Even	v_2	Real	Even
$A\phi_o^2$	Real	Even	y_2	Real	Even
$A\phi_e\phi_o$	Real	Odd	iz	Imaginary	Odd

purely real or imaginary with the same symmetries; we list the corresponding terms in Table 1.

Thus, all the introduced variables in Table 1 are purely real. The quantities a , v , and y denote the Fourier transforms of the pupil function, and even- and odd-WF aberrations, respectively:

$$a = \mathcal{F}\{A\}, \tag{5}$$

$$v = \mathcal{F}\{A\phi_e\}, \tag{6}$$

$$y = \text{Im}\{\mathcal{F}\{A\phi_o\}\}. \tag{7}$$

Using the definitions, the second-order PSF approximation can be written as

$$p = |a + iv - y - 0.5v_2 - 0.5y_2 - iz|^2, \tag{8}$$

which simplifies to

$$p = a^2 + v^2 + y^2 - 2ay + \xi, \tag{9}$$

where the first four terms constitute the first-order approximation—in terms of WF expansion—and the second-order component is

$$\xi = 0.25v_2^2 + 0.25y_2^2 + z^2 - av_2 - ay_2 + 0.5v_2y_2 + yv_2 + yy_2 - 2vz. \tag{10}$$

The above equations are best illustrated by an example. Consider a purely sinusoidal WF having a peak-to-valley value of 1.0 rad and an rms error of 0.37 rad; alternative examples can be seen, for instance, in [17]. The WF and the resulting PSF image are shown in Fig. 2. The WF causes two main sidelobes and more sidelobes with significantly lower intensity; one pair is shown in Fig. 2.

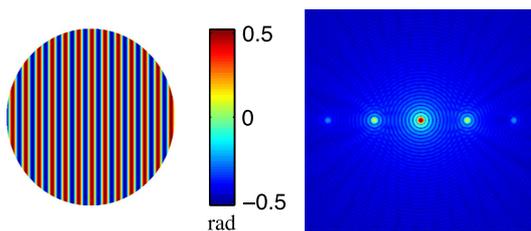


Fig. 2. Left, a purely sinusoidal WF. Right, resulting image raised to the power of 0.2 to compress the dynamic range.

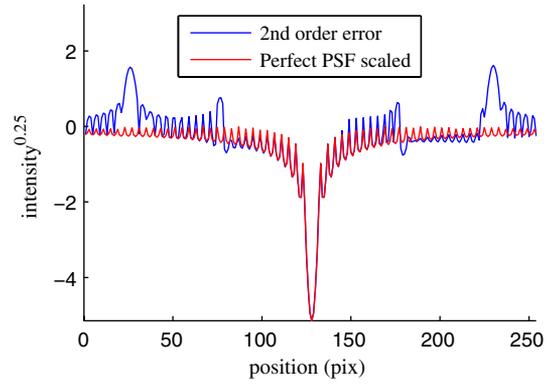


Fig. 3. Radial cuts of the second-order component ξ , defined in Eq. (10), and an inverted and scaled perfect PSF, a^2 .

Figure 3 shows a radial cut of the second-order component ξ for the example WF. Its most significant terms are av_2 and ay_2 , and therefore the perfect image (a^2) scaled by a negative coefficient approximates ξ reasonably well. This term is responsible for the energy conservation by reducing the Strehl ratio [13]. The first-order approximation always has a Strehl ratio of 1.

Thus, an improved first-order approximation can be obtained by subtracting a scaled version of a^2 from the first-order PSF approximation; the scaling coefficient needs to be adjusted such that the maxima of the perfect PSF and the approximation are the same. The radial cuts of the PSF approximations are illustrated in Fig. 4. The improved first-order approximation captures the main lobe and the first pair of sidelobes quite well, but the secondary sidelobes are missed.

However, for a WF with an rms error of less than 1 rad, the improved first-order approximation is often sufficient, and it can be formulated as

$$p = a^2 + y^2 + v^2 - 2ay - \left(1 - \frac{\max(p_n)}{\max(a^2)}\right)a^2, \tag{11}$$

where p_n denotes the recorded image normalized to the same energy as the perfect PSF:

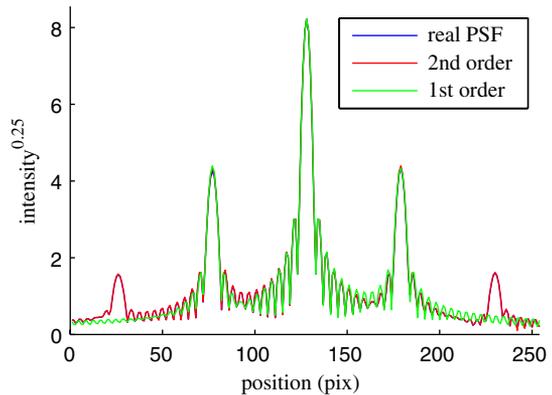


Fig. 4. Radial cuts of the perfect PSF, its improved first-order approximation, and the second-order approximation. The latter is virtually identical to the perfect PSF.

$$p_n = p_m \frac{\sum_{x,y} a^2(x,y)}{\sum_{x,y} p_m(x,y)}, \quad (12)$$

where (x,y) denotes the detector pixel coordinates and p_m is the raw image. Therefore, to simplify the notations, it is convenient to define a modified normalization of a PSF:

$$p' = p_n + \left(1 - \frac{\max(p_n)}{\max(a^2)}\right) a^2, \quad (13)$$

where the normalized image, p' , has the same maximum as a^2 .

To solve the WF using Eq. (11), we follow the procedure of [16], which is repeated here for convenience.

The recorded image is normalized and broken into its even and odd parts. It then holds that

$$p'_e = a^2 + v^2 + y^2, \quad (14)$$

$$p'_o = 2ay. \quad (15)$$

The odd component of the WF is then easily reconstructed by first solving y using Eq. (15) and then using the inverse of Eq. (7). Due to noise and approximation errors, the direct application of Eq. (15), however, would result in division by excessively small values. We compensate for this by using a regularization as in [16]:

$$y = \frac{ap'_o}{2a^2 + \epsilon}, \quad (16)$$

where ϵ is a small number. We found it best to set ϵ to a value of 50–500 times the measured noise level of the recorded images.

To compute the even-WF component, we need additional information in the form of phase diversity. We assume that a second, previously recorded image is known, and it was obtained with a known phase change compared to p . The even component of its normalized version can be written as

$$p'_{e2} = a^2 + (v + v_d)^2 + (y + y_d)^2, \quad (17)$$

where v_d and y_d are the even and odd Fourier components of the phase diversity, obtained in analogy to Eqs. (6) and (7).

Using Eqs. (14) and (17), we can solve v (the even-phase component in Fourier space) and write it as

$$v_s = \frac{p'_e - p'_{e2} - v_d^2 - y_d^2 - 2yy_d}{2v_d}. \quad (18)$$

However, this formula is highly sensitive to noise due to the subtraction of two very similar images. Therefore, as also seen in [16], we use Eq. (18) only to

compute the signs of v . A more robust form follows from the use of Eq. (14):

$$v = \text{sign}(v_s) |p'_e - a^2 - y^2|^{0.5}, \quad (19)$$

where we use the absolute value to avoid taking the square root of negative values that occur due to noise and approximation errors; this was observed to work better than zeroing the negative values. The even-WF component is then computed in the same way as the odd one, by using Eq. (19) and the inverse of Eq. (6).

B. Practical Aspects

To use the FF algorithm as presented here, it is necessary to have a WF correcting device, a deformable mirror or SLM with a known phase response. It is then possible to translate the desired phase change to appropriate WF-corrector-command signals. Appropriate mapping can be created using the standard adaptive-optics-calibration procedures as in [14] or, as we do here, with the help of a differential optical transfer function (dOTF)-based calibration method [15]. The method is based on determining the SLM phase (and transmission) response when the control signal is changed in different pixel blocks. This data is then used to find an affine transform that maps the location of each SLM pixel to its physical location in the pupil plane.

We also assume that the collected images are sufficiently sampled: without aberrations the full width at half maximum of the PSF has to be at least two pixels. If the detector is undersampled, aliasing prevents using the intensity images as described in Section 2.A. Large oversampling is also not desired since it increases the computational requirements.

The phase array, ϕ , needs to be sampled with sufficient resolution to also model the pupil aperture, A , with good accuracy. The values we use (170×170) are sufficient for our purpose; we expect no significant sampling errors when implementing Eqs. (6) and (7) as fast Fourier transforms (FFTs). However, we need to zero-pad the recorded images such that the FFTs correctly implement the Fourier transforms in Eqs. (5), (6), and (7); the sampling of the arrays a , v , and y need to match the pixels of the camera. The amount of zero padding is determined by the sampling coefficient

$$q = \frac{N_{\text{arr}}}{N_{\text{pup}}}, \quad (20)$$

where N_{arr} is the dimension of the FFT array and N_{pup} is the size of ϕ . We use the dOTF method as discussed in [15] to find q . The method is based on the use of localized phase diversity at the pupil border, which makes it possible to very straightforwardly create an array where the pupil shape can be directly seen. The parameter q is calculated by comparing the sizes of the pupil and the dOTF array.

When performing the FFT to obtain the phase from v and y , we combine the two real-valued FFTs to a single complex FFT [13]:

$$A\phi = \mathcal{F}^{-1}\{w(v + iy)\}, \quad (21)$$

where w is a windowing function. It implements filtering necessary for numerical regularization; typically, high-spatial frequencies are detected with higher uncertainty and need to be damped to obtain feasible reconstructions. The regularization is also needed with noiseless images because the weak-phase solution provides only approximate WFs. In this work, we have used a concave parabola, whose width can be adjusted depending on the noise level. An optimum filter is the subject of future studies.

To implement the iterative feedback loop to optimize the WF error, we use a standard, leaky-integrator control. The WF-corrector shape at time step k is calculated as

$$\theta_k = g_l\theta_{k-1} - gA\phi_{k-1}, \quad (22)$$

where g_l is the leaky gain; θ_{k-1} is the previous WF-corrector shape; g is the integrator gain; and $A\phi_{k-1}$ is the most recent small-phase solution, computed using the two most recent images using Eq. (21).

The integrator gain, g , determines the trade-off between convergence speed and stability; a small gain results in slow convergence, while a high gain means the image noise causes larger errors after the algorithm has converged. Excessively small gain would also make the use of phase-diversity information difficult.

The leaky gain is another regularization parameter. A value of $g_l = 1$ would be equal to a standard integrator, and it would be optimal in the case of no errors, with the equation $p = |\mathcal{F}\{A \exp(i\phi)\}|^2$ perfectly describing the system. Values $g_l < 1$ introduce WF aberrations at every time step, preventing the system from reaching a perfect state. However, that also prevents creeping instabilities from destroying the performance. The result is a stable convergence at a level with a slightly higher residual-WF error.

3. Fast & Furious Gerchberg–Saxton

The obvious limitation of the FF algorithm is the assumption of the pupil amplitudes being even. This holds reasonably well for most of the optical systems having a circular shape, possibly with a central obstruction. However, to achieve the optimal focal-plane-WF sensing with a high-order system not suffering from other limiting factors, it is necessary to consider imaging models where the pupil amplitudes can have an arbitrary shape.

We have approached the problem by combining the FF-style weak-phase solution and a version of the GS algorithm. The new algorithm is referred to as FF-GS in the following.

As with the GS algorithm, we maintain an iteratively updated estimate of the unknown quantities, in our case, the pupil amplitudes. The pupil amplitude estimate, phase diversities, and the recorded images are used to calculate the focal-plane field; it requires three Fourier transforms and the use of a weak-phase approximation. Next, a Fourier transform is used to propagate the field to the pupil plane. The propagation results in improved estimates for the pupil-plane amplitudes and the WF. The schematic illustration of the FF-GS algorithm is shown in Fig. 5.

The FF-GS computation procedure forms a loop that could be iterated several times to obtain improved WF estimates. However, we found that in practice it is sufficient to run only two iterations before applying the WF correction with the obtained estimate. As with FF, the WF correction yields another image and phase-diversity information, which are used to compute the following correction step.

Next, Section 3.A describes the algebra that we use to compute the focal-plane electric field during the FF-GS procedure. Then, Section 3.B explains the details of the iterative computation, and Section 3.C discusses practical issues we face when implementing the algorithm.

A. More General Weak-Phase Solution

In this section, we assume that an approximation of the pupil amplitudes (denoted here as A) is known; as a first step, a top-hat function is sufficient in the case of an unobstructed, round pupil. The estimates are

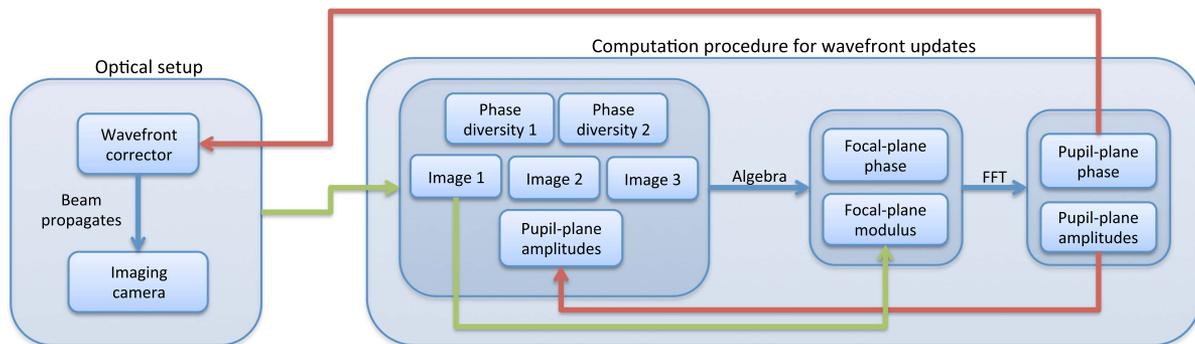


Fig. 5. Schematic illustration of the FF-GS algorithm.

updated iteratively, and we will make no restrictive assumptions about A .

We assume that three images are collected and that the corresponding phase-diversity information is known. The images are normalized according to Eq. (13), and it holds approximately that

$$p'_1 = |e_1|^2 = |\mathcal{F}\{A + iA(\phi)\}|^2, \quad (23)$$

$$p'_2 = |e_2|^2 = |\mathcal{F}\{A + iA(\phi + \phi_{d1})\}|^2, \quad (24)$$

$$p'_3 = |e_3|^2 = |\mathcal{F}\{A + iA(\phi + \phi_{d2})\}|^2, \quad (25)$$

where e_1, e_2 , and e_3 are the electric fields corresponding to the images; ϕ is the unknown pupil-plane phase; and ϕ_{d1} and ϕ_{d2} are the known phase diversities applied to successively recorded images.

When counting the number of unknown variables, one can see that it might be possible to solve the unknown phase using only two images with Eqs. (23) and (24). However, we found the following procedure with three images to be better. In addition to making the algebra easier, it is also significantly more robust since more information is available to compensate the errors in the estimate of A . Using even more images could potentially still improve the results, but studying this is outside the scope of this paper.

Instead of solving the phase directly, we use phase-diversity information to find the electric field at the focal plane. The electric field corresponding to Eq. (23) can be written as

$$e_1 = (a_r + \alpha) + i(a_i + \beta), \quad (26)$$

where

$$a_r = \text{Re}\{\mathcal{F}\{A\}\},$$

$$a_i = \text{Im}\{\mathcal{F}\{A\}\},$$

$$\alpha = -\text{Im}\{\mathcal{F}\{A\phi\}\},$$

$$\beta = \text{Re}\{\mathcal{F}\{A\phi\}\}.$$

The unknown coefficients α and β can be found by solving the equations that follow when subtracting Eq. (23) from Eqs. (24) and (25). The subtraction cancels all the nonlinear terms and results in linear equations

$$\begin{bmatrix} 2\alpha_{d1} & 2\beta_{d1} \\ 2\alpha_{d2} & 2\beta_{d2} \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} c_1 \\ c_2 \end{bmatrix}, \quad (27)$$

where

$$\alpha_{d1} = -\text{Im}\{\mathcal{F}\{A\phi_{d1}\}\},$$

$$\beta_{d1} = \text{Re}\{\mathcal{F}\{A\phi_{d1}\}\},$$

$$\alpha_{d2} = -\text{Im}\{\mathcal{F}\{A\phi_{d2}\}\},$$

$$\beta_{d2} = \text{Re}\{\mathcal{F}\{A\phi_{d2}\}\},$$

and

$$\begin{aligned} c_1 &= p'_2 - p'_1 - (2a_r\alpha_{d1} + 2a_i\beta_{d1} + \alpha_{d1}^2 + \beta_{d1}^2), \\ c_2 &= p'_3 - p'_1 - (2a_r\alpha_{d2} + 2a_i\beta_{d2} + \alpha_{d2}^2 + \beta_{d2}^2). \end{aligned} \quad (28)$$

We solve the coefficients α and β by inverting the 2×2 matrix in Eq. (27). The matrix has full rank if the used phase diversities are linearly independent. We found this generally to be the case when applying the algorithm, and therefore it was unnecessary to use any regularization methods. The coefficients can then be substituted into Eq. (26) to compute the focal-plane electric field. However, this estimate would again be very prone to noise due to the subtraction of similar images, as shown in Eq. (28). Therefore, it is better to use the directly measured modulus and use only the phase information following from Eq. (26). This then gives a more robust focal-plane estimate:

$$e_1 = |p'_1|^{0.5} \exp[i \arg((a_r + \alpha) + i(a_i + \beta))]. \quad (29)$$

The following section explains the details of how this is then combined with the GS approach.

B. Iterative Computation Procedure

As the previous section indicates, we first record three images. The phase diversity can be chosen freely, as long as its peak-to-valley stays below 1 rad. We use the FF algorithm at the initial steps.

Then, using the collected data, we perform computations to calculate a new WF update. The WF update is applied, and another image with different phase-diversity information is collected. The three most recent images are then used again to calculate the next phase correction to be applied. We continue until the algorithm converges.

The computation consists of a cycle of two successive GS-like iterations. The complete process consists of the following steps:

1. Take the pupil amplitudes, A , estimated at the previous iteration. Use the procedure in Section 3.A to calculate the focal-plane electric field corresponding to p_2 , the second most recent image. This is done by solving α and β in Eq. (27) and using formula

$$e_2 = |p'_2|^{0.5} \exp[i \arg((a_r + \alpha) + i(a_i + \beta))].$$

Here, the images could be rearranged appropriately: p_2 should be the reference and the phase diversities interpreted accordingly. However, we found $\arg(e_2) \approx \arg(e_1)$ to be a sufficient approximation.

2. Compute the pupil-plane electric field corresponding to the image p_2 . This is done by Fourier transforming the focal-plane field:

$$E_2 = \mathcal{F}^{-1}\{e_2\}.$$

3. Update the current estimate of the pupil amplitudes:

$$A = |E_2|.$$

4. With the new pupil-amplitude estimate, repeat the procedure in Section 3.A to compute the electric field for image p_1 , the most recent image.

5. Compute the pupil-plane field corresponding to image p_1 :

$$E_1 = \mathcal{F}^{-1}\{e_1\}.$$

6. Calculate the final phase estimates for the phase and pupil amplitudes:

$$\phi = \arg(E_1), \quad (30)$$

$$A = |E_1|. \quad (31)$$

The estimates of ϕ are then used in the feedback loop in the same way as with the FF algorithm.

C. Practical Aspects

The issues faced in practice by FF-GS implementation differ slightly from the simple FF.

Since the pupil amplitudes are not constrained, the imaging model is potentially much more accurate. In practice, indeed, we found that it was not necessary to apply any windowing filters to dampen the high-spatial frequencies in the WFs reconstructed with FF-GS. The normal feedback loop, as described by Eq. (22), was sufficient regularization for the optimal performance.

It was also not necessary to introduce any *ad hoc* restrictions to constrain the pupil amplitudes. The values obtained from Eq. (31), at any time step, do have a significant deviation from the actual pupil amplitudes, but this appears to be a minor issue for the convergence of the algorithm. Moreover, averaging the values of A over several iterations produces nonbiased results.

However, the heavier reliance on the phase-diversity information makes the algorithm more prone to stability issues. To increase the stability, we found it helpful to introduce other *ad hoc* techniques.

In the feedback loop, we apply amplitude gains. Just as formulated in Eq. (22), we multiply the applied phase correction obtained from Eq. (30) by the estimated amplitudes. This helps to prevent abrupt phase changes at points where $|E_1|$ has a very small value; at those points, the determination of the complex phase is likely to fail. In fact, we also set ϕ to zero at points where $|E_1| < 0.3$. This reduces the speed of convergence, but has no impact on the accuracy of the converged solution.

Finally, additional regularization is used in case of numerical issues when the algorithm has converged. We observed that occasionally, every 10th iteration or so, the FF-GS algorithm produces wildly incorrect results. This is related to the fact that the solution of Eq. (27) requires phase-diversity information. Once

the applied phase corrections become very small, the corresponding diversity information becomes unreliable.

To make sure that such violent phase changes will not cause troubles, we simply restrict the magnitude of the applied phase change. If the rms value of the change exceeds the mean of ten previous changes, we scale it down to the mean value.

4. Hardware Used

To test the algorithms, we created a simple setup that consists of one SLM and an imaging camera. The former is a reflective device (BNS P512) having a screen of 512×512 pixels, a fill factor of 83.4%, and a pixel pitch of $15 \mu\text{m} \times 15 \mu\text{m}$. The SLM is able to create a phase change of 2π rad at the used wavelength, and its control signal is coded with 6 bits.

The imaging camera is a Basler piA640-210gm, which has a resolution of 648×488 pixels and a dynamic range of 12 bits. As a light source, we use a fiber-coupled laser diode (Qphotonics' QFLD-660-2S) having a wavelength of 656 nm.

A schematic figure of the setup is shown in Fig. 6. The beam goes first through a diaphragm, and it is then collimated such that it hits an area of 245×245 pixels on the SLM. The device reflects several subbeams due to strong diffraction effects, and we use only the zeroth-order beam; it is directly imaged onto the camera (beam numerical aperture = 0.037). The other subbeams cause no adverse effects. Before and after the SLM, we place two linear polarizers that are rotated such that their orientation matches the one of the SLM.

The SLM phase and transmittance responses are measured with the dOTF method described in [15]. The resulting measurements are shown in Fig. 7. The maximum control voltage causes $\sim 2\pi$ phase shift at 656 nm.

The used SLM couples the transmittance and phase change; the transmittance gradually increases when a larger phase shift is introduced with the SLM. For phase changes of less than 1 rad, the transmittance is $\sim 25\%$ lower compared to what is seen when a change of more than ~ 4 rad is introduced.

To create a mapping between the pupil-plane coordinates and the SLM pixels, we again use the dOTF

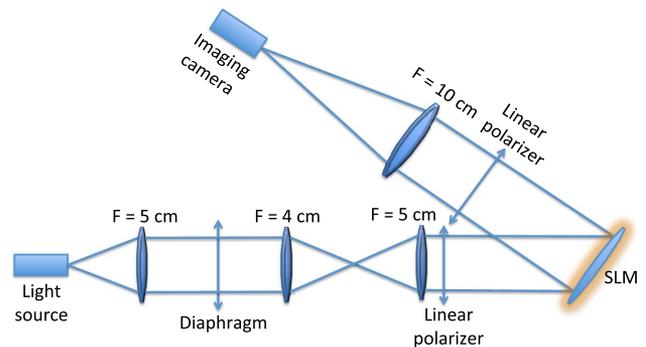


Fig. 6. Schematic view of the used hardware. The lenses are standard 1 inch doublets. The beam diameter is 3.7 mm at the SLM.

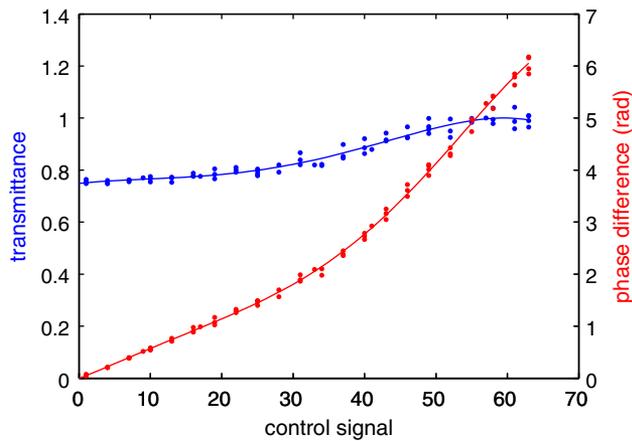


Fig. 7. SLM phase and amplitude responses. The dots indicate individual measurements. The lines show fifth-order polynomial fits to the data.

method and affine transforms as described in [15]. This time, however, we make the dOTF record the best focus to avoid issues with the nontelecentric beam. To compensate for signal-to-noise problems, we take more images to average out the noise; it takes ~ 2 h to create one dOTF array. This makes the process also more vulnerable to internal turbulence in the setup; the recorded images are blu0072red such that the low-spatial frequencies in the images become distorted, and we have to mask out the center of the obtained dOTF arrays.

Figure 8 shows the modulus of the best-focus dOTF array recorded with the whole SLM at zero-control voltage. Although the center of the array is masked, it is still perfectly usable for the calibration process of [15], and we can accurately determine the PSF sampling as defined by Eq. (20): $q = 3.76 \pm 0.01$.

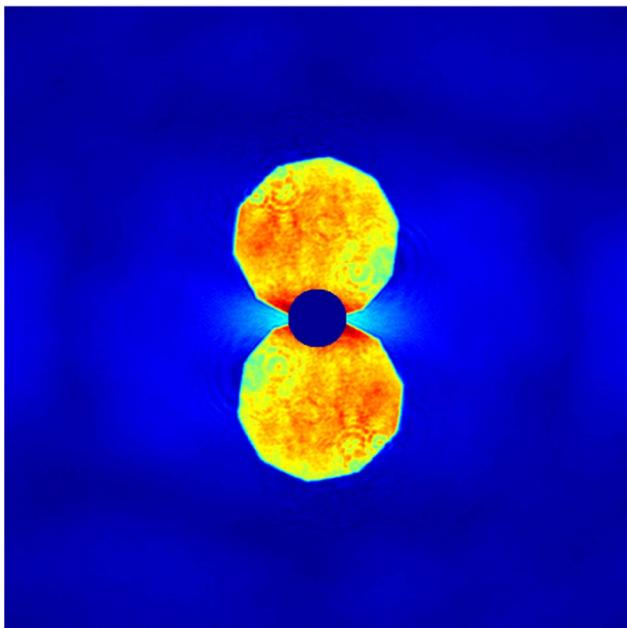


Fig. 8. Modulus of an averaged dOTF array.

The resulting SLM calibration is valid as long as the position of the SLM stays fixed with respect to the imaging camera, and the phase-response of the device does not change. In our setup, we found this to be the case for at least one month—from the initial calibration to the last measurements reported in this paper.

As discussed in [15], the resolution of the controlled phase is a free parameter when calculating the affine mapping for the SLM calibration. We obtained good results when using $\sim 30\%$ fewer pixels than are actually used by the SLM. Thus, we selected the size of the controlled phase array as $N_{\text{pup}} = 170$. The resulting FFT-array dimension is then $N_{\text{arr}} = 640$.

When recording images for the FF and FF-GS algorithms, we use the same high-dynamic range (HDR) imaging approach as in [15]. Several snapshot images are taken with different exposure times, and we combine the images to extend the dynamic range and compensate noise. Each single-exposure component in one HDR image is an average over 40–200 images, and we used in total 16 exposure times (2, 5, 12, 25, 50, 100, 200, 400, 750, 1100, 1450, 1800, 2150, 2500, 2850, and 3200 ms). It took ~ 15 s to record one HDR image. Increasing the integration even further does not significantly improve the performance of the WF correction algorithms.

Although the imaging camera has a resolution of 640×480 pixels, we use a smaller area for convenience reasons. After acquiring the image, we crop an array of 320×320 pixels such that the PSF maximum is in the center. Outside of the region, we did not observe any significant amount of light.

To detect all the spatial frequencies corrected by the controlled-phase array of 170×170 pixels, however, we would need an array of 640×640 pixels. Thus, it is possible that our control algorithms introduce high-spatial frequencies that scatter light outside of the observed image. However, with FF, this is mitigated by the applied low-pass filter. With FF-GS, we observed no stability issues with the high-spatial frequencies, although no explicit regularization measures were taken.

5. Results

This section illustrates the results of the FF and FF-GS algorithms. We consider only a single case: the WF to be corrected is what the camera sees at the beginning, when no voltage is applied to the SLM. We call this the initial situation.

We concentrate on the ultimate accuracy that the algorithms can achieve in a low-noise regime. Our earlier publication [14] describes in more detail the FF performance in the presence of more noise. We showed that the algorithm works, but only the lower spatial frequencies can be reconstructed. Now, we study a case that is typical for a high-order adaptive-optics test bench, and the noise level is chosen so that FF-GS offers an advantage over FF. With higher noise, FF is more robust.

Section 5.A illustrates the properties of the converged algorithms as measured with our test setup. Section 5.B shows a more detailed comparison of the measurements and simulations with the actual hardware modeled in sufficient detail. Finally, Section 5.C presents a simulation-based error budget that quantifies the effects of different error sources.

A. Performance of the Algorithms

For the results shown here, we have optimized the free parameters (FF regularization coefficient ϵ , the width of the FF filtering window w , leaky gain g_l , and loop gain g) such that the converged WF quality is best; the convergence speed has lower priority.

The width of the filtering window used by the FF algorithm was chosen to be 320×320 pixels, the same as the recorded images. However, during the first 10 iterations, we used a narrower window (width of 80 pixels) to avoid introducing errors at the high-spatial frequencies. After the lower spatial frequencies are corrected, it is safe to increase the window size.

The optimal values for feedback-loop gains were $g = 0.3$, $g_l = 0.97$ (with FF), or $g_l = 0.999$ (with FF-GS), and ϵ was 250 times the determined noise level in the images.

For the FF algorithm, we also need to determine the pupil amplitudes, A . We use a perfect top-hat function having a size of $N_{\text{pup}} \times N_{\text{pup}}$, where the choice of N_{pup} is explained in Section 4. It might be possible to improve the results by adjusting A based on the actual pupil shape, but this is outside the scope of this paper.

With these settings, both FF and FF-GS converge in 20–50 iterations to a situation where the Strehl ratio has increased from $\sim 75\%$ to $\sim 99\%$ (a more detailed analysis can be found in Section 5.B). After the convergence, the control law, Eq. (22), gives phase updates that are negligible compared to the shape of the WF corrector, θ_k . However, we run the algorithm for a total 400 iterations to make sure that no creeping instabilities occur.

Figure 9 illustrates the typical WFs we obtained after the convergence. Due to the applied low-pass filter, FF yields WFs smoother than FF-GS; otherwise, however, they match well. The repeatability of the experiments appears reasonable: the converged WF shapes have experiment-to-experiment differences of at most ~ 0.2 – 0.3 rad. The spread of the FF-GS results tends to be smaller compared to FF, and we see that also the higher spatial frequencies are produced in a repeatable way.

Figure 10 shows the reconstructed pupil amplitudes. The top left shows an average of A following the application of Eq. (31) during a total of 400 FF-GS iterations with phase updates. It can be compared with the dOTF modulus shown next to it, and we see that the shape of the diaphragm and several bigger dust particles are correctly recovered. However, it is obvious that all the finer details are lost, and the very lowest spatial frequencies also deviate

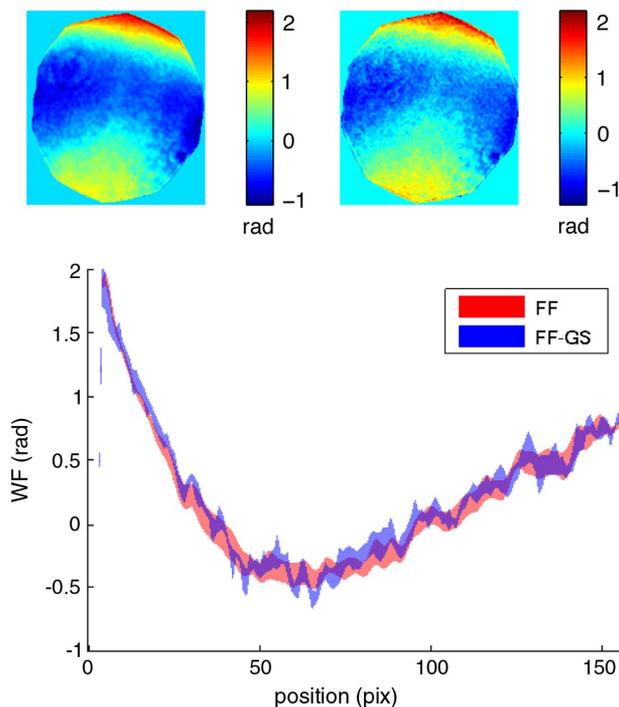


Fig. 9. Top row, typical WF shapes (170×170 pixels) of the SLM after the convergence of FF and FF-GS. Bottom, radial cuts through the WFs; the shaded area shows the range (minima and maxima) of five independent measurements.

from each other. The plot at the bottom of Fig. 10 shows radial cuts of five similarly obtained pupil amplitudes, and we see that all the features in the

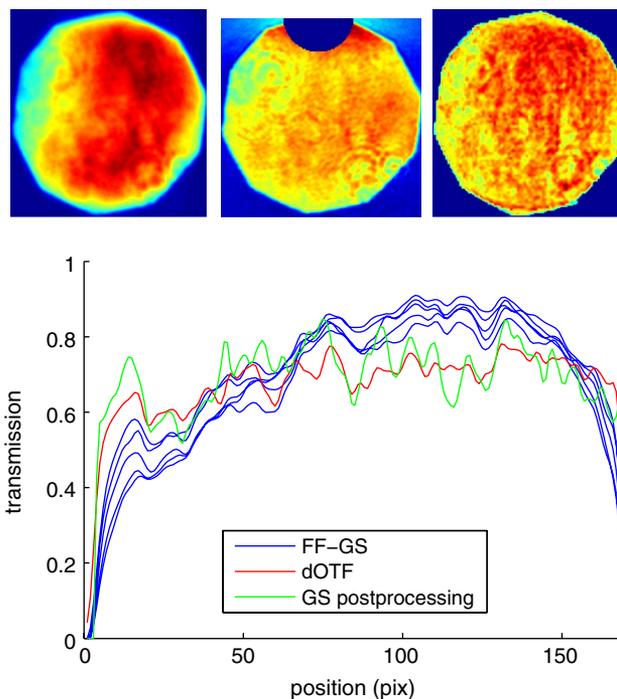


Fig. 10. Top row, pupil amplitudes (170×170 pixels) reconstructed with different methods. Left, FF-GS. Middle, dOTF (same as in Fig. 8). Right, GS postprocessing from a converged PSF. Bottom, radial cuts through the pupil amplitudes; five independent measurement runs are shown for FF-GS.

pupil amplitudes are nevertheless repeatedly reconstructed in the same way.

To obtain an improved reconstruction of the finer details in the pupil amplitudes, we use the PSF that results after the FF-GS algorithm has converged. We assume that all the remaining speckles are caused by the amplitude aberrations and reconstruct—with a Gerchberg–Saxton-style algorithm—a pupil that would create such a pattern. This is shown in the upper right of Fig. 10, and we can see that it much better matches the dOTF reconstruction in Fig. 8. Later, we use this pattern in simulations for analysis purposes.

The differences between the independent measurement series shown here are a combination of actual small changes in the hardware and uncertainty caused by noise and systematic errors. It is difficult to separate those two effects, and therefore we continue the analysis with the help of numerical simulations.

B. Comparison of Measurements and Simulations

To simulate the optical setup, we assume that the algorithms correct WFs shown in Fig. 9 with pupil amplitudes similar to what is shown in Fig. 10. We created three case studies reflecting the variability in the converged results.

In the simulations, we consider eight different sources of errors that needs to be modeled explicitly. They are:

1. SLM quantification. We use only 6 bits to control the WF. The plots shown in Fig. 7 are used to round the simulated WF correction to what would happen in practice.

2. PSF sampling. The WF and the resulting PSF are sampled internally by a factor of two higher than what the hardware controls or observes. The control algorithms use rebinned PSFs, and the simulated-WF correction is interpolated bilinearly from the reconstruction at a resolution of 170×170 .

3. Image noise and dynamic range. We estimate the read-out noise of the HDR images to be at a level of $2.2 \cdot 10^{-6}$ of the image maximum. Gaussian random noise is added to the simulated PSFs. The HDR images have maximum values $\sim 4 \cdot 10^8$, corresponding to about 29 bits, and this is also modeled in the simulations.

4. Background level. Standard background subtraction is performed on the PSF images, but a small error will still remain. Therefore, we add a constant background level, $2.7 \cdot 10^{-6}$ of the image maximum, to the simulated PSFs.

5. Nonperfect pupil. Instead of the perfect top-hat function, we use pupil amplitudes similar to what is illustrated in the top right of Fig. 10.

6. Amplitude aberrations. We simulate the coupling of the WF and the transmission of the SLM as illustrated by Fig. 7.

7. Alignment errors. Although the dOTF calibration is rather accurate, some error could still be present in the affine transform that we use to map

the WF to the SLM pixels. The simulations indicate that if the transform has a mismatch corresponding to a rotation larger than 0.4° , FF and FF-GS would be unstable. In practice, with the used hardware, we saw no hints these of instabilities. Therefore, a rotation error of 0.4° represents the maximum misregistration that the WF control algorithms are likely to experience.

8. Tip-tilt error. Internal turbulence in the optical setup causes frame-to-frame WF variations, which can be approximated to a degree as small shifts of the recorded images. We measured the difference of the center-of-gravity between two consecutive PSFs recorded with the HDR method, and it was found to be on average 0.025 pixels. This error cannot be taken into account by the phase-diversity approach, and we model its impact on the performance.

Figure 11 shows the remaining WF error as a function of time step. The simulation plots show the exact error, but the measured value is estimated from the data. Here, we have estimated the rms error from the corresponding PSF images only. At first, we estimated the Strehl ratios using the method seven in [18], and the result was converted to an rms error using the expression $S = \exp(-\sigma^2)$. The resulting estimates are highly sensitive to the estimation of the pupil amplitudes, which we know only approximately (Fig. 10). Thus, the y axis in the lower plot in Fig. 11 is not directly comparable to the simulation

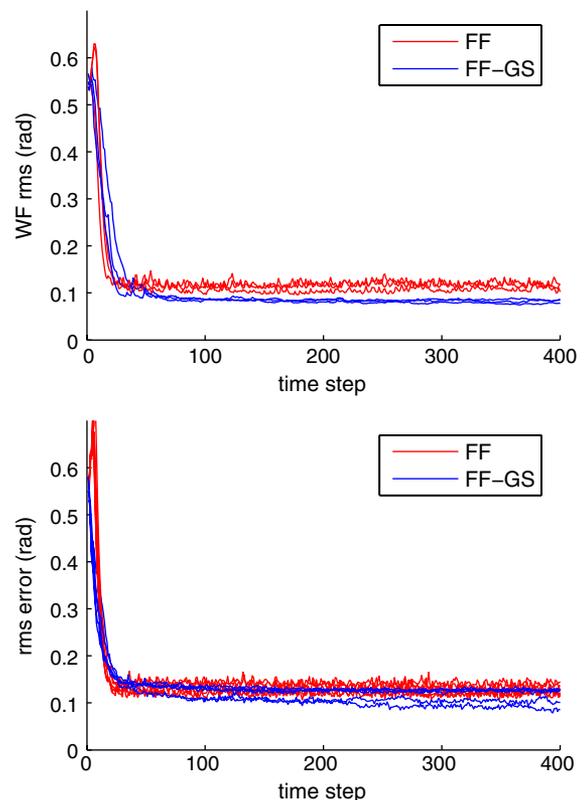


Fig. 11. Tip-tilt-removed residual WF error as a function of time step. Top, simulations (real value). Bottom, measurements (estimation from PSF images).

plot; alternative estimates that are more easily compared are shown later in this section.

Nevertheless, the speed of the convergence is clearly seen. Both FF and FF-GS reduce the WF rms error from ~ 0.5 rad rms to ~ 0.1 in ~ 50 iterations. FF converges about 50% faster, but it is plagued by the overshoot at the beginning; it would require an adaptive optimization of the low-pass filter to properly handle it.

Regarding the simulations, it is obvious that the FF-GS improves the performance over FF: the rms error is 0.08 rad as compared to 0.12 rad. This is largely due to the smaller value of the leaky integrator gain that we had to apply to make the FF stable.

Regarding the measurements, we can see a similar pattern, but we also see that the FF-GS has two modes: the estimate of the residual rms error is either ~ 0.10 rad or ~ 0.13 rad. The modes are related to the finite sampling of the CCD detector. Our models do not explicitly constrain the position of the PSF at the detector, which means that a random subpixel tip-tilt component—different in each independent measurement series—is left in the residual WF. The algorithms converge to a state that remains stable, but the different remaining tip-tilt components can cause significant changes in the measured maximum intensity, which affects our Strehl-ratio-estimation process. When inspecting the recentered PSFs carefully, as shown later in this section, no significant differences between the PSFs can be seen.

A more detailed investigation reveals that the convergence of the WF correction depends on the spatial frequency; low-frequency features are reconstructed faster. Figure 12 illustrates this by showing how an average intensity in different regions of the field changes as a function of time step. We show three different regions representing low, medium, and high-spatial frequencies; the locations correspond to Airy rings 2–4, 12–17, and rings further than 30. Since we consider only small WF aberrations, the shown intensity values are directly proportional to the average-power-spectral density at the matching frequency bands.

Both simulations and measurements show a similar pattern, although the absolute levels are higher in simulations due to differences in noise. At low-spatial frequencies, both FF and FF-GS peak at iterations 5–10. FF converges in total in ~ 20 iterations, and FF-GS takes ~ 20 iterations more, although some cases show intensity reduction even until ~ 100 iterations. At medium-spatial frequencies, the peak occurs at approximately iteration 15, and the algorithms need in total ~ 30 iterations to reach an intensity level $\sim 6\%$ lower than at the beginning. FF saturates at that level, but 30 additional iterations with FF-GS reduce the intensity in total $\sim 15\%$ from the initial level. At high-spatial frequencies, FF requires almost 50 iterations to converge to a level 15% lower than the initial intensity (in simulations, the reduction is only a few percentages due to higher

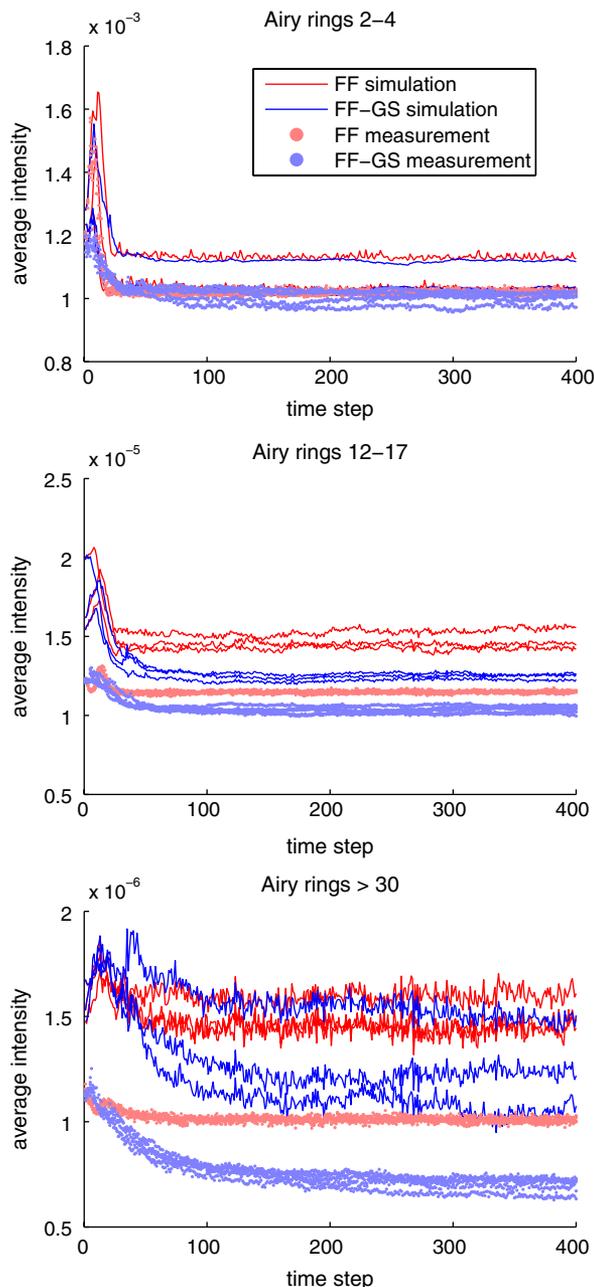


Fig. 12. Average intensity at different parts of the field. Three cases are shown: the field corresponding to Airy rings 2–4, Airy rings 12–17, and Airy further than 30.

noise). FF-GS, on the other hand, converges faster than FF, but still 150 iterations are needed to reduce the intensity $\sim 35\%$. The measurements show marginally better intensity reduction, but that requires almost 300 iterations.

The residual-WF error can obviously also be estimated using the control data that the algorithms themselves provide through Eqs. (21) and (30); the corresponding results are shown in Fig. 13.

The first striking feature is that the simulations and the measurements produce practically identical patterns. After the convergence, the WF estimates of the FF algorithm have an rms error of 0.12–0.18 rad

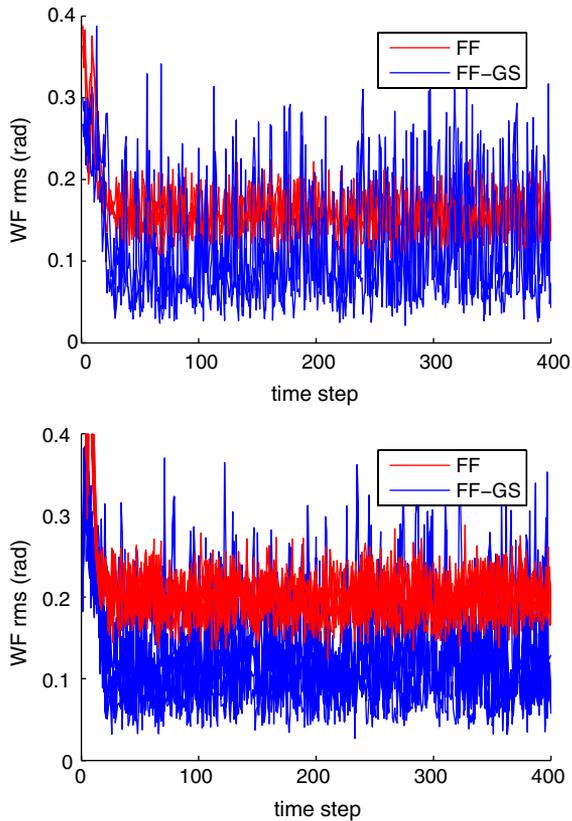


Fig. 13. Residual-WF error as a function of time step. Values calculated from the actual estimates used by the algorithms. Top, simulations. Bottom, measurements.

in the simulations and 0.15–0.24 rad in the measurements. There appears to be no obvious structure in how the error varies between consecutive iterations. Since the actual correction is an average over several consecutive measurements, the actual remaining WF error can be smaller than the instantaneous estimates of 0.12–0.24 rad. In the simulations, the error was observed to be ~ 0.12 rad, and we have no reason to assume the situation with the actual hardware would be different; our estimate for the remaining WF rms error is ~ 0.15 rad.

With the FF-GS algorithm, the issue is slightly more complicated since some of the WF estimates fail when the algorithm approaches the optimum. The reason for this—the phase-diversity failure—is discussed in Section 3.C. This is seen as prominent spikes in the plots in Fig. 13, although most of the rms error values are concentrated around 0.1 rad. In the simulations, the actual rms error of the residual WF is ~ 0.08 rad, and a similar value is seen in the actual measurements.

Four examples of the actual PSF images are shown in Fig. 14:

- A. the initial PSF (measured when the SLM pixels are set to zero),
- B. the simulated perfect PSF resulting from the pupil amplitudes shown in Fig. 10,
- C. simulated PSF after the convergence of the FF-GS algorithm,

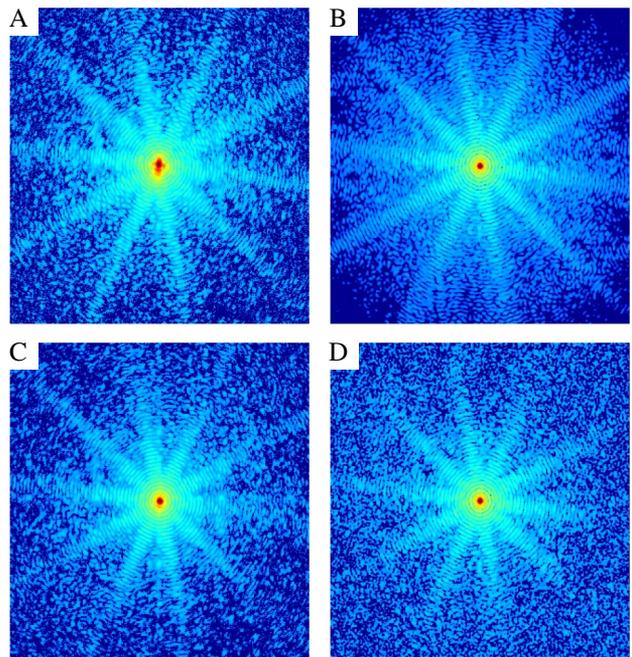


Fig. 14. Examples of PSF images (320×320 pixels) raised to the 0.1 power. A, initial, measured; B, perfect, simulated; C, converged FF-GS, measured; D, converged FF-GS, simulated.

D. measured PSF after the convergence of the FF-GS algorithm.

All the PSFs have a similar star pattern with 10 radial beams gradually fading towards the edges of the images. These are caused by the blades of the diaphragm, whose shape is shown in Figs. 8 and 10.

The initial PSF corresponds to a WF like in Fig. 9: a clearly deformed core, but still easily recognizable Airy rings 3–20.

The simulated, noiseless, and aberration-free PSF shows the speckles that we expect to remain due to the nonflat pupil amplitudes. The dust, dirt, and inhomogeneities of the SLM create a significant transmission distortion dominated by high-spatial frequencies. This causes the halo of irregularities on top of the pattern of the perfect diffraction rings. In addition, we can see a few stronger speckles and speckle groups at a distance of approximately Airy rings 12–18. These can be attributed to the larger dust particles also clearly visible in the FF-GS estimated pupil amplitudes in Fig. 10.

When comparing the measured and simulated PSFs after the FF-GS algorithm has converged, we find no significant differences. Both PSFs have a regular core, which appears to match exactly the perfect PSF up to the fourth diffraction ring. At least 26 diffraction rings are at least partially visible. A comparison with the perfect PSF shows that several strong speckles can be identified in all the images, but the halo after the fourteenth diffraction ring outside the star-like beams, close to the detection limit of the camera, is dominated by speckles with no obvious structure.

A more detailed comparison can be obtained by inspecting the radially averaged profiles of the PSFs. Before taking the radial average, we shift, using Fourier transforms, the PSFs to have the center of gravity at the location of the perfect PSF. The results are shown in Fig. 15.

The profiles show that both the FF and FF-GS algorithms, in both the simulated and measured cases, converge to a situation very close to the perfect simulated PSF; no significant differences are seen up to the first 13 (simulated) or 20 (measured) diffraction rings. After this, we can see that the performance of both algorithms slowly deviates from the perfect PSF, the intensity being a factor of ~ 5 (simulated) or $\sim 2\text{--}3$ (measured) higher at borders. At the distances corresponding to diffraction rings 20 and higher, FF-GS is typically $\sim 20\text{--}30\%$ better in reducing the intensity as compared to FF.

In total, we can recognize at least 30 diffraction rings before the speckle noise makes the PSF structure too blurry to observe any structure. Nevertheless, compared to the initial PSF, both algorithms reduce the intensity of scattered light throughout the whole used field. However, in the simulated case, the difference is not significant after the 34th diffraction ring. In the measured case, on the other hand, the light intensity is reduced by a

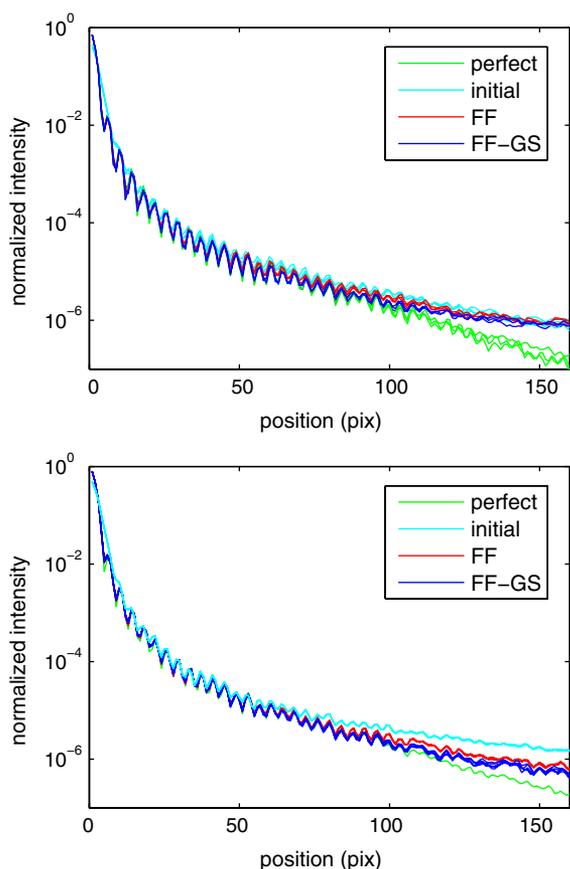


Fig. 15. Averaged radial profiles of PSF images. Upper, simulated, with three study cases shown. Lower, measured, with results from five independent runs shown. The perfect PSF is identical to the one in the upper plot.

factor of $\sim 2\text{--}3$ also at the edge of the recorded image. This difference between the simulations and measurements is due to a combined effect of differences in actual noise levels, WFs, and pupil transmission.

C. Error Budget

Finally, we show an error budget that illustrates the impact of the different error sources in the optical setup.

In the ideal case, we have no noise and a perfectly circular pupil that is—in the case of FF—exactly known. The perfect case also uses exactly the same imaging model in both the WF reconstruction and when simulating the PSF images: a zero-padded FFT with a WF modeled at a resolution of 170×170 pixels.

We sequentially simulate each of the error sources listed in Section 5.B. The resulting rms errors in the converged WFs are listed in Table 2.

In theory, both algorithms should reach zero WF error in the perfect case. However, in the case of FF, we still have to use numerical regularization to maintain stability, and this compromises the performance in the error-free case. This could be improved by optimizing the codes, but it is not done here; the codes are optimized for the performance with all the error sources present.

The most severe error source for the FF algorithm, as expected, is the amplitude aberrations: instead of the ideal rms error of 0.03 rad, we are limited to an error of 0.11 rad. Similar errors are also seen if the imaging model does not exactly match the actual hardware; this was tested when simulating the WF and PSF with double sampling (case 2 in Table 2). The double sampling was also used in the misalignment simulation. The different error sources are coupled, so they do not add up quadratically. In the presence of all the error sources, we end up having a residual WF error of ~ 0.12 rad.

With the FF-GS algorithm, we can radically reduce the problems of the unknown pupil aberrations. The transmission we used in simulations, however, had significant fluctuations, creating speckles similar to what the WF aberrations do. Therefore, the WF reconstruction problem is difficult to make unambiguous, and we saw a small residual rms error of 0.02 rad.

Table 2. Error Budget

	FF ^a	FF-GS ^a
0. No errors	0.03 ± 0.01	0.00 ± 0.00
1. SLM quantification	0.04 ± 0.01	0.02 ± 0.00
2. PSF sampling 2×	0.08 ± 0.01	0.01 ± 0.00
3. Image noise	0.05 ± 0.01	0.05 ± 0.00
4. Background level	0.04 ± 0.01	0.01 ± 0.00
5. Nonperfect pupil	0.11 ± 0.00	0.02 ± 0.01
6. Amplitude aberrations	0.12 ± 0.01	0.04 ± 0.01
7. Alignment errors	0.08 ± 0.01	0.01 ± 0.00
8. TT instability	0.03 ± 0.01	0.04 ± 0.01
9. All errors	0.12 ± 0.01	0.08 ± 0.00

^aThe residual WF rms errors (rad) at spatial frequencies falling within the used images.

The FF-GS is limited by the combined effect of read-out noise (0.05 rad), TT instability (0.04 rad) and the fact that the SLM couples the phase and transmission changes (0.04 rad). All the error sources add up quadratically, which indicates that they are largely independent.

When comparing the FF and FF-GS, we see that a significant improvement can be obtained with the FF-GS algorithm; the residual WF rms error is reduced from 0.12 to 0.08 rad. However, the method is more sensitive to uncertainties and noise: the tip-tilt jitter in our hardware has no influence on the FF while being a major error source in the FF-GS algorithm.

6. Conclusions and Discussion

We have demonstrated the performance of two numerically efficient focal-plane WF-sensing algorithms: the Fast & Furious and its extension, the Fast & Furious Gerchberg–Saxton algorithm.

Both algorithms do an excellent job in calibrating static aberrations in an adaptive- or active-optics system: we demonstrated an increase in the Strehl ratio from ~ 0.75 to 0.98 – 0.99 with our optical setup.

Although the FF-GS algorithm is more prone to noise, we observed a clear improvement. With our hardware—a high-resolution spatial-light modulator as the WF corrector—we estimate the remaining residual WF rms error to be ~ 0.15 rad with FF and ~ 0.10 rad with FF-GS. The difference occurs mostly at spatial frequencies corresponding to the twentieth and further Airy rings.

Simulations with error sources comparable to our hardware show very similar results. This increases our confidence that the estimated performance indicators are reliable, and the simulated error budget also confirms the unknown amplitude aberrations as the main limitation of the FF algorithm in the considered framework.

To our knowledge, this is the first time that such focal-plane-sensing methods have been demonstrated with $\sim 30,000$ deg of freedom—and in the case of FF-GS, with twice the number of free parameters to estimate the pupil amplitudes.

The sampling at the detector was such that the controlled WF of 170×170 pixels would have been enough to correct all spatial frequencies inside an image of 640×640 pixels. However, as we recorded only an image of 320×320 pixels, we had no direct observations of the higher controlled spatial frequencies. Simulations indicate that this resulted in a small amount of light being scattered outside the recorded field, but this amount was too small to be easily detected in our optical setup.

We put no particular effort into optimizing the codes; all the software was implemented in Matlab, and it was run on a standard Windows PC. Still, the required computation time was negligible compared to the ~ 15 s we needed to collect data for a single HDR image. We implemented the FF algorithm with two 640×640 FFTs per iteration step

(one FFT transferring the phase-diversity information into the focal plane could likely be replaced by a convolution, as explained in [13]). Our FF-GS implementation used eight FFTs per iteration, and that could also potentially be optimized.

As with all focal-plane WF sensing techniques, the algorithms work best if a monochromatic light source is available. With a chromatic light source having a sufficiently small bandwidth, perhaps $\sim 10\%$, the algorithms would still work, but only with a limited corrected field. With special chromatic optics (such as in [19]) or an integral field unit, it may be possible to use the algorithms with even wider bandwidth.

Currently, we have only demonstrated a case where an unobstructed PSF is detected, and the WF is driven to be flat. To make the algorithms more interesting for astronomical applications in extreme adaptive optics or ultra-high contrast-imaging, a few extensions would be necessary.

First, we should consider how coronagraphs and diffraction suppression optics will affect the techniques. In practice, this would mean that the core of the PSF would not be detected, and we would need to also consider the moduli in a part of the focal-plane field as free parameters.

Second, instead of flattening the WF, we should optimize the contrast at a certain part of the field. This would mean calculating a WF shape that, in the same way as in [9,20,21], minimizes the light in certain regions of the field at the cost of increasing it in other parts; the updated algorithm should then drive the WF to this desired shape. A similar problem is faced if phase plates are used to create diffraction suppression, for instance as in [22]. Also in such a case, it is necessary to drive the WF to a particular shape that is far from flat.

Another potentially interesting application is a real-time application, for instance as a high-order, second-stage sensor in an adaptive-optics system. The computational load is manageable, and a successful system would greatly simplify the hardware design compared to a conventional adaptive optics approach. However, issues such as the requirement for small aberrations, chromaticity, temporal lag in the phase diversity, and the limited dynamic range of the camera—and therefore photon noise—are major challenges.

References

1. R. W. Gerchberg and W. O. Saxton, "A practical algorithm for the determination of the phase from image and diffraction plane pictures," *Optik* **35**, 237–246 (1972).
2. J. R. Fienup, "Phase retrieval algorithms: a comparison," *Appl. Opt.* **21**, 2758–2769 (1982).
3. J. J. Green, D. C. Redding, S. B. Shaklan, and S. A. Basinger, "Extreme wave front sensing accuracy for the Eclipse coronagraphic space telescope," *Proc. SPIE* **4860**, 266–276 (2003).
4. R. S. Burruss, E. Serabyn, D. P. Mawet, J. E. Roberts, J. P. Hickey, K. Rykoski, S. Bikkannavar, and J. R. Crepp, "Demonstration of on sky contrast improvement using the modified Gerchberg-Saxton algorithm at the Palomar Observatory," *Proc. SPIE* **7736**, 77365X (2010).
5. J. Sauvage, T. Fusco, G. Rousset, and C. Petit, "Calibration and precompensation of noncommon path aberrations for

- extreme adaptive optics,” *J. Opt. Soc. Am. A* **24**, 2334–2346 (2007).
6. P. Riaud, D. Mawet, and A. Magette, “Nijboer-Zernike phase retrieval for high contrast imaging. Principle, on-sky demonstration with NACO, and perspectives in vector vortex coronagraphy,” *Astron. Astrophys.* **545**, A150 (2012).
 7. B. Paul, L. M. Mugnier, J.-F. Sauvage, M. Ferrari, and K. Dohlen, “High-order myopic coronagraphic phase diversity (COFFEE) for wave-front control in high-contrast imaging systems,” *Opt. Express* **21**, 31751–31768 (2013).
 8. C. Vérinaud, M. Kasper, J.-L. Beuzit, R. G. Gratton, D. Mesa, E. Aller-Carpentier, E. Fedrigo, L. Abe, P. Baudoz, A. Boccaletti, M. Bonavita, K. Dohlen, N. Hubin, F. Kerber, V. Korkiakoski, J. Antichi, P. Martinez, P. Rabou, R. Roelfsema, H. M. Schmid, N. Thatte, G. Salter, M. Tecza, L. Venema, H. Hanenburg, R. Jager, N. Yaitskova, O. Preis, M. Orecchia, and E. Stadler, “System study of EPICS: the exoplanets imager for the E-ELT,” *Proc. SPIE* **7736**, 77361N (2010).
 9. A. Giv'oni, R. Belikov, S. Shaklan, and J. Kasdin, “Closed loop, DM diversity-based, wavefront correction algorithm for high contrast imaging systems,” *Opt. Express* **15**, 12338–12343 (2007).
 10. C. S. Smith, R. Marinică, A. J. den Dekker, M. Verhaegen, V. Korkiakoski, C. U. Keller, and N. Doelman, “Iterative linear focal-plane wavefront correction,” *J. Opt. Soc. Am. A* **30**, 2002–2011 (2013).
 11. S. Meimon, T. Fusco, and L. M. Mugnier, “Lift: a focal-plane wavefront sensor for real-time low-order sensing on faint sources,” *Opt. Lett.* **35**, 3036–3038 (2010).
 12. F. Martinache, “The asymmetric pupil Fourier wavefront sensor,” *Publ. Astron. Soc. Pac.* **125**, 422–430 (2013).
 13. C. U. Keller, V. Korkiakoski, N. Doelman, R. Fraanje, R. Andrei, and M. Verhaegen, “Extremely fast focal-plane wavefront sensing for extreme adaptive optics,” *Proc. SPIE* **8447**, 844721 (2012).
 14. V. Korkiakoski, C. U. Keller, N. Doelman, R. Fraanje, R. Andrei, and M. Verhaegen, “Experimental validation of optimization concepts for focal-plane image processing with adaptive optics,” *Proc. SPIE* **8447**, 84475Z (2012).
 15. V. Korkiakoski, N. Doelman, J. Codona, M. Kenworthy, G. Otten, and C. U. Keller, “Calibrating a high-resolution wavefront corrector with a static focal-plane camera,” *Appl. Opt.* **52**, 7554–7563 (2013).
 16. R. A. Gonsalves, “Small-phase solution to the phase-retrieval problem,” *Opt. Lett.* **26**, 684–685 (2001).
 17. M. D. Perrin, A. Sivaramakrishnan, R. B. Makidon, B. R. Oppenheimer, and J. R. Graham, “The structure of high Strehl ratio point-spread functions,” *Astrophys. J.* **596**, 702–712 (2003).
 18. L. C. Roberts, Jr., M. D. Perrin, F. Marchis, A. Sivaramakrishnan, R. B. Makidon, J. C. Christou, B. A. Macintosh, L. A. Poyneer, M. A. van Dam, and M. Troy, “Is that really your Strehl ratio?” *Proc. SPIE* **5490**, 504–515 (2004).
 19. O. Guyon, “High sensitivity wavefront sensing with a nonlinear curvature wavefront sensor,” *Publ. Astron. Soc. Pac.* **122**, 49–62 (2010).
 20. F. Malbet, J. W. Yu, and M. Shao, “High-dynamic-range imaging using a deformable mirror for space coronagraphy,” *Publ. Astron. Soc. Pac.* **107**, 386–398 (1995).
 21. P. J. Bordé and W. A. Traub, “High-contrast imaging from space: speckle nulling in a low-aberration regime,” *Astrophys. J.* **638**, 488–498 (2006).
 22. J. L. Codona and R. Angel, “Imaging extrasolar planets by stellar halo suppression in separately corrected color bands,” *Astrophys. J., Lett.* **604**, L117–L120 (2004).