

PROCEEDINGS OF SPIE

[SPIDigitalLibrary.org/conference-proceedings-of-spie](https://spiedigitallibrary.org/conference-proceedings-of-spie)

Speckle nulling wavefront control for Palomar and Keck

Michael Bottom, Bruno Femenia, Elsa Huby, Dimitri Mawet, Richard Dekany, et al.

Michael Bottom, Bruno Femenia, Elsa Huby, Dimitri Mawet, Richard Dekany, Jennifer Milburn, Eugene Serabyn, "Speckle nulling wavefront control for Palomar and Keck," Proc. SPIE 9909, Adaptive Optics Systems V, 990955 (27 July 2016); doi: 10.1117/12.2233025

SPIE.

Event: SPIE Astronomical Telescopes + Instrumentation, 2016, Edinburgh, United Kingdom

Speckle nulling wavefront control for Palomar and Keck

Michael Bottom,^a Bruno Femenia,^b Elsa Huby,^c Dimitri Mawet,^a Richard Dekany,^a Jennifer Milburn,^a Eugene Serabyn^d

^aCalifornia Institute of Technology, Pasadena, CA, United States;

^bW.M. Keck Observatory, Waimea, HI, United States;

^cUniversity of Liège, Liège, Wallonia, Belgium;

^dJet Propulsion Lab, Pasadena, CA, United States

ABSTRACT

We present a speckle nulling code currently being used for high contrast imaging at the Palomar and Keck telescopes. The code can operate in open and closed loop and is self-calibrating, requiring no system model and minimal hand-coded parameters. Written in a modular fashion, it is straightforward to port to different instruments. It has been used with systems operating in the optical through thermal infrared, and can deliver nearly an order of magnitude improvement in raw contrast. We will be releasing this code to the public in the near future.

Keywords: High contrast imaging, coronagraphy, adaptive optics, wavefront correction

1. INTRODUCTION

The largest current barrier to imaging extrasolar planets is the presence of “speckle” aberrations, which arise from errors in figure or transmissiveness of optical elements after the wavefront sensor of the adaptive optics system. These speckles show up in the focal plane as bright points of light that look similar to the point-spread function, typically tens to thousands of times brighter than any planetary companions that might be present in the image. To a large degree, observing approaches on ground-based telescopes are oriented towards removing speckles in post-processing. This includes techniques such as angular differential imaging,¹ spectral differential imaging,² and reference star differential imaging.

Rather than only remove speckles in post processing, it is possible to remove them “optically.” This has the advantage that the fundamental photon shot noise level is reduced, so all post-processing algorithms will have better performance, as they are all eventually limited by shot noise. (Post-processing techniques, of course, are completely compatible with optical speckle reduction.) There are a number of different techniques to reduce speckles optically. The most powerful general approach is called “electric field conjugation” (EFC),³ where the deformable mirror response function (in electric field) is characterized and modeled, small perturbations are put on it to determine the phase over the focal plane; then the linearized response function is used to compute the corrections in a specified control region. EFC has been used in a number of stable high contrast imaging testbeds, with contrasts of better than 10^{-8} being achieved in some cases.⁴

Electric field conjugation has not found wide implementation on high contrast imaging instruments on telescopes, perhaps because it requires an accurate model of the imaging system to achieve good performance. In many instruments, system states can change quite rapidly due to changing gravity vectors, temperature, or other states. This motivates us to consider a different algorithm for robust performance on non-ideal systems. An alternative to EFC is speckle nulling, an algorithm which identifies bright points of light in the image, uses the deformable mirror to create an artificial “speckle” at the same location, and calculates the phase of the original speckle by changing the phase of the artificial speckle.^{5,6} Finally, it uses the deformable mirror to put a speckle of the opposite phase on the original offending speckle, which makes it disappear. Doing this to many speckles at a specified location in the image plane will create a region of lowered intensity and higher contrast.

Send correspondence to M.B. at mbottom@caltech.edu

In this work, we will first present the background principles of speckles; how they arise, how they propagate, and how it is possible to (partially) correct them. Then we will describe the main points behind speckle nulling, and discuss our specific algorithmic implementation. Finally, we will present results of our speckle nulling code on different instruments at the Palomar and Keck observatories.

2. BACKGROUND PRINCIPLES

2.1 Speckles in high contrast imaging

Speckles are formed by optical errors in the imaging system, and can be decomposed into two different types; “phase-type” and “amplitude-type”. A phase-type speckle is formed by a phase error on an optical surface, such as an imperfect figure, while an amplitude-type speckle is formed by variations in reflectivity, possibly caused by an imperfect polish. Significantly, phase-type errors can turn into amplitude-type errors and vice versa, via the Talbot effect, with the conversion taking place over distances of approximately $2a^2/\lambda$, where a is the spatial period of the error and λ is the wavelength; this ratio typically corresponds to about 1 meter of optical propagation. Speckles in the focal plane generally consist of both phase-type and amplitude-type, and a single speckle may be a superposition of both types. Here, we follow the explanation in Traub and Oppenheimer,⁷ mostly conserving their notation. Readers are encouraged to consult that source for a thorough and insightful exposition of these phenomena.

The electric field of a pure plane wave at a pupil of the imaging system can be written as

$$E(r, t) = A_0 e^{i(k \cdot r - 2\pi f t + \phi_0)} \quad (1)$$

In the following discussion, since we are concerned with the final intensity, not electric field, we ignore the quickly oscillating terms of type $e^{i2\pi f t}$, where f is the frequency of light, on the order of 100 THz for these systems. Similarly, we restrict the coordinates to a single variable, x , which corresponds to the coordinate at a one-dimensional pupil that runs from $-D/2$ to $D/2$, D being the diameter of the pupil. Let us now consider static phase and amplitude type errors at the pupil, so that the amplitude and phase of the wave is not constant, but has a dependence on the position in the pupil

$$A(x) = a(x) e^{i\phi_p(x)} \quad (2)$$

Here, we let $a(x) = A_0(1 - \phi_a(x)) \approx e^{\phi_a(x)}$ be a constant electric field amplitude plus some small perturbation in the pupil, and let $\phi_p(x)$ be any phase aberrations at the same location. (All quickly oscillating $e^{i2\pi f t}$ terms are dropped). We then absorb the amplitude aberrations and phase aberrations into a single complex phase error $\phi(x)$ so that

$$A(x) = A_0 e^{i\phi(x)} \quad (3)$$

$$\phi(x) = \phi_p(x) + i\phi_a(x) \quad (4)$$

Consider the case where the error consists of a single phase error* (such as a ripple in a mirror) of amplitude a , spatial frequency x_0 and “pupil phase” α . By pupil phase, we refer to the pupil decentration of the perturbation— with the convention that a pure cosine with the peak at the center of the pupil has pupil phase 0. The single phase ripple here is described by

$$\phi(x) = a \cos(2\pi x/x_0 + \alpha) \quad (5)$$

To see the effect of this perturbation in the image plane (where the detector is located), we recall there is a Fourier transform between the electric field of the pupil and image plane in an optical system. (Even if a detector measures intensity and not electric field, it is the field that carries the phase information explicitly).

*More realistic errors can be decomposed into linear superpositions of such errors via Fourier analysis.

Therefore, the electric field at the following image plane will be given by the Fourier transform of this ripple ϕ , as follows:

$$A_{\text{image}}(\theta) = \int_{-D/2}^{D/2} e^{i\phi(x)} e^{i2\pi\theta x/\lambda} dx \quad (6)$$

$$\approx \int_{-D/2}^{D/2} (1 + ia \cos(2\pi x/x_0 + \alpha)) e^{i2\pi\theta x/\lambda} dx \quad (7)$$

$$= \int_{-D/2}^{D/2} e^{i2\pi\theta x/\lambda} dx + \frac{ia}{2} \int_{-D/2}^{D/2} e^{i(2\pi\theta x/\lambda + \alpha + 2\pi x/x_0)} dx + \frac{ia}{2} \int_{-D/2}^{D/2} e^{i(2\pi\theta x/\lambda - \alpha - 2\pi x/x_0)} dx \quad (8)$$

where the approximation is justified by the phase ripple being a small error. The last line follows from the second by expanding $\cos x = (e^{ix} + e^{-ix})/2$. We also note that the dx here refers to the pupil coordinate.

Letting the Fourier transform of a function f be denoted by $\mathcal{F}[f(x)] = \hat{f}(\theta)$, we now recall the following Fourier transform relations

$$\mathcal{F}[e^{i2\pi gx} f(x)] = \hat{f}(\theta - g) \quad (9)$$

$$\mathcal{F}[f(x/c)] = c\hat{f}(c\theta) \quad (10)$$

$$\mathcal{F}[\text{rect}(x)] = \frac{\sin(\pi\theta)}{\pi\theta} \quad (11)$$

$$(12)$$

where

$$\text{rect}(x) = \begin{cases} 0 & |x| > \frac{1}{2} \\ \frac{1}{2} & |x| = \frac{1}{2} \\ 1 & |x| < \frac{1}{2} \end{cases}$$

Here we use the variable θ to denote angle. If a ripple in the pupil plane has a spatial frequency of cycles/cm, or cm^{-1} , then θ should strictly be in units of cm, just as the Fourier transform of a time series has units of t^{-1} , or Hz. However, since an imaging system essentially converts angles into positions, we casually write the position coordinate on the detector as an angle.

The electric field in the focal plane may then be written in the following form

$$A(\theta) = \int_{-\infty}^{\infty} \text{rect}\left(\frac{x}{D}\right) e^{i2\pi\theta x/\lambda} dx \quad (13)$$

$$+ \frac{iae^{i\alpha}}{2} \int_{-\infty}^{\infty} \text{rect}\left(\frac{x}{D}\right) e^{i(2\pi\theta x/\lambda + 2\pi x/x_0)} dx + \frac{iae^{-i\alpha}}{2} \int_{-\infty}^{\infty} \text{rect}\left(\frac{x}{D}\right) e^{i(2\pi\theta x/\lambda - 2\pi x/x_0)} dx \quad (14)$$

Defining

$$A_0(\theta) = \frac{D \sin(\pi\theta D/\lambda)}{\pi\theta D/\lambda} \quad (15)$$

where $A_0(\theta)$ is recognized as a sinc function, or the 1-D Fourier transform of the pupil, (this would be a 2-d Airy function in the case of a circular aperture). Finally, carrying out the integrals using the Fourier relations above, at the following image plane, the electric field will be described by

$$A(\theta) = A_0(\theta) + \frac{1}{2} a e^{i(\pi/2 - \alpha)} A_0(\theta - \lambda/x_0) + \frac{1}{2} a e^{i(\pi/2 + \alpha)} A_0(\theta + \lambda/x_0) \quad (16)$$

(Note that constants of the form i and $e^{i\alpha}$ are absorbed into the complex exponential so that phases, or complex phasors, are consistently represented.) The previous result should be interpreted as a central PSF (an Airy function) plus two speckles of diminished amplitude at positions $\pm\lambda/x_0$, the same shape as the central PSF. Note that the “pupil phase” α determines the electric field phases, but that each speckle has a different electric field phase, with the speckle on one side having a phase of $\pi/2 + \alpha$, and the other having a phase of $\pi/2 - \alpha$. In this case, a deformable mirror correction of $\phi(x) = a \cos(2\pi x/x_0 + \alpha + \pi)$ will completely cancel the speckle, as can be immediately seen from Equation 5.

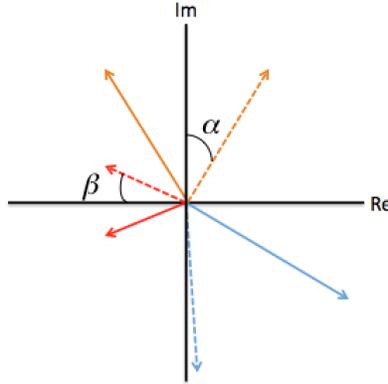


Figure 1: Complex plane representation of speckle phasors. A speckle at positions λ/x_0 containing phase-type (orange arrows) and amplitude-type (red arrows) contributions with electric field phases $\pi/2 \pm \alpha$ and $\pi \pm \beta$, respectively, where α and β are the pupil phases. The dashed arrows refer to the components at $-\lambda/x_0$, and the solid arrows refer to the components at λ/x_0 . A speckle at λ/x_0 may be nulled out by the electric field phasor corresponding to the blue solid arrow (as can be seen by computing the phasor sums when lying the arrows head to tail). This will not simultaneously null out the speckle at $-\lambda/x_0$, whose null solution is shown by the blue dashed arrow.

Similarly, an amplitude-type error across the pupil *at the same spatial frequency*, but different amplitude b and pupil phase β will be described by

$$\phi(x) = ib \cos(2\pi x/x_0 + \beta) \tag{17}$$

with the corresponding intensity in the focal plane being given by

$$A(\theta) = A_0(\theta) + \frac{1}{2}be^{i(\pi-\beta)}A_0(\theta - \lambda/x_0) + \frac{1}{2}be^{i(\pi+\beta)}A_0(\theta + \lambda/x_0) \tag{18}$$

Note that the speckle shows up at the same focal plane locations regardless of the pupil phase, ($\theta \pm \lambda/x_0$), but the electric field phases depend on the pupil phase in a different way, being $\pi \pm \beta$ in the case of the amplitude ripple rather than $\pi/2 \pm \alpha$, as in the case of the phase ripple. This is important, because a two-sided speckle with combined phase and amplitude components cannot be canceled by the deformable mirror, which can only provide a phase correction. However, two-sided speckle may be nulled on one side by the deformable mirror. See Figure 1 for a visual explanation of this.

In order to null out a speckle, it is necessary to know the phase of its electric field at that point. Consider a speckle with both phase-type and amplitude-type contributions, as before. This speckle can be described by a single electric field phase and single amplitude, ie, a single complex number—this is simply the static component of the electric field of the light at that point. Explicitly combining the first phase and first amplitude focal plane terms, we can write the speckle as:

$$A(\theta) = \frac{1}{2} \left[ae^{i(\pi/2-\alpha)} + be^{i(\pi-\beta)} \right] A_0(\theta - \lambda/x_0) \quad (19)$$

$$= re^{i\phi} A_0(\theta - \lambda/x_0) \quad (20)$$

$$\phi = \tan^{-1} \left[\frac{a \sin \alpha - b \cos \beta}{a \cos \alpha + b \sin \beta} \right] \quad (21)$$

$$r = \frac{1}{4} \sqrt{a^2 + b^2 + 2ab \sin(\alpha - \beta)} \quad (22)$$

2.2 Principles of speckle nulling

The previous section described how speckles arise, how they can be due to both phase-type and amplitude-type errors, and how they can be described by a single electric field phase and amplitude at a particular spot in the focal plane. Furthermore, it was shown how speckles can be corrected with the deformable mirror on only one side of the focal plane, or partially corrected on both sides. In this section, an algorithm to do that will be outlined.

While the previous section was completely general, in this section we will describe a specific technique to remove speckles. Speckle nulling is just one technique among many to measure and cancel the electric field in the focal plane. In some sense, it is the “simplest” way to solve this problem, compared to more advanced methods like electric field conjugation.³ This simplicity has both benefits and drawbacks. The advantages are that there is very little system-specific knowledge that is needed, and it is robust to changing instrumental conditions and non-idealities. The main disadvantage is that the convergence rate is slower than methods that take into consideration explicit instrumental parameters, so for a fixed number of iterations, the contrast gain will be somewhat lower.

In this section, we will describe the phase sensing and correction part of the speckle nulling algorithm. It turns out that the sensing part of the algorithm is essentially identical to phase-shifting interferometry; where a “reference” beam of light interferes with light reflected off some test surface, the reference beam is phase-shifted, and the pixel-to-pixel interferogram variations encode the phase. The main difference is that rather than a test surface, light will be interfering in the focal plane.

Suppose a speckle in the focal plane has an electric field of

$$A(\theta) = a_s e^{i\phi_s} A_0(\theta - \lambda/x_0) \quad (23)$$

where a_s is the speckle amplitude and ϕ_s is the phase, assuming A_0 has been appropriately normalized. Suppose the deformable mirror overlays a speckle on it of known electric field phase ϕ_{DM} , and known amplitude a_{DM} . The interference of these two waves will lead to the following intensity pattern in the focal plane, where the intensity is given by the squared sum of the electric fields:

$$I = |a_s e^{i\phi_s} + a_{DM} e^{i\phi_{DM}}|^2 \quad (24)$$

$$= a_s^2 + a_{DM}^2 + 2a_s a_{DM} \cos[\phi_s - \phi_{DM}] \quad (25)$$

$$= I_s + I_{DM} + 2\sqrt{I_s I_{DM}} \cos[\phi_s - \phi_{DM}] \quad (26)$$

The equation above has omitted quickly oscillating terms of type $e^{2\pi i f t}$ which do not show up in the final answer, and also removed the spatial and structural dependence given by $A(\theta)$. Furthermore, this equation is assumed to hold at pixels corresponding to a particular speckle location, where the deformable mirror precisely overlays its own “artificial” speckle. By choosing $\phi_{DM} = [0, \pi/2, \pi, 3\pi/2]$, it is possible to have the following intensities corresponding to each of the phases:

$$I_1 = I_s + I_{DM} + 2\sqrt{I_s I_{DM}} \cos[\phi_s] \quad (27)$$

$$I_2 = I_s + I_{DM} - 2\sqrt{I_s I_{DM}} \sin[\phi_s] \quad (28)$$

$$I_3 = I_s + I_{DM} - 2\sqrt{I_s I_{DM}} \cos[\phi_s] \quad (29)$$

$$I_4 = I_s + I_{DM} + 2\sqrt{I_s I_{DM}} \sin[\phi_s] \quad (30)$$

To solve for the speckle phase, then, it is possible to simply compute

$$\phi_s = \tan^{-1} \left[\frac{I_4 - I_2}{I_1 - I_3} \right] \quad (31)$$

Another quantity of interest is the speckle visibilities $\gamma(x, y)$, equal to the average modulation of the speckle divided by the mean value

$$\gamma(x, y) = \frac{2\sqrt{I_s I_{DM}}}{I_s + I_{DM}} = \frac{2\sqrt{(I_4 - I_2)^2 + (I_1 - I_3)^2}}{I_1 + I_2 + I_3 + I_4} \quad (32)$$

It should be noted that a large visibility (of near 1) is optimal, as it will lead to the highest fringe contrasts and hence highest phase accuracy. The maximal visibility occurs when the artificial and real speckles are the same amplitude; this can be formally proven, but is intuitively obvious when considering the extreme case of an artificial speckle so dim that there is no detectable modulation ($\gamma = 0$). Here, this situation is equivalent to not performing the experiment; one would be as well off just guessing the electric field phase.

From the above discussion, the speckle amplitude may be computed too, as a_{DM} is known, since it is controlled by the experimenter. However, this is not done in this case, because the preceding discussion made it clear that in order to measure phase, it is necessary to have a sufficient visibility, and hence have I_s and I_{DM} be approximately the same. To match the electric field amplitudes when overlaying the artificial speckles, one must have an a priori idea of how intensity in the focal plane (measured in counts, photoelectrons, etc) depends on amplitude (in nanometers or microns) on the deformable mirror. This is precisely the same problem, so solving for a_s is not really necessary. We will discuss how a_s is determined below.

3. OVERVIEW OF CODE

The speckle nulling code is written in Python 2.7. The code consists of four kinds of files: the configuration files (instrument and nulling), the instrument modules, the calibration programs, and the nulling programs. The instrument modules are specific to the particular system (camera and adaptive optics system) being used, and contain the basic commands to take images, change the shape of the deformable mirror, etc. The instrument calibration files control the basic parameters of the instruments, like default exposure times of the science camera, communication ports, and other details. The calibration programs determine the optical parameters of the system, like deformable mirror to camera rotation angle, pixel pitch in units of λ/D , and how the deformable mirror amplitude translates to image intensity, and write these parameter to the nulling configuration file. Finally, the nulling configuration file controls how many iterations to run, how many speckles to null at once, and other

In normal operations, the only files that ever need to be modified are the configuration files, which are human-readable and do not require any programming knowledge. Of course, the instrument modules will need to be written for different cameras or adaptive optics systems, but many observatories already have APIs for their instruments. All that may be required is a simple set of Python “wrapping” functions to execute the commands through the observatory API, then return and open any image files in a common directory.

3.1 Calibration

3.1.1 Dark frames, flat fields, and bad pixel corrections

The presence of “hot”, dead, or hyperactive pixels in the image plane can cause slower convergence in the speckle control algorithm. It was found that operating on cleaned images can give better performance than operating on the raw frames, though this depends on the relative quality of the science camera in question. Regardless, the first step in the speckle nulling code is to prompt the user to set the system up in a flatfield, dark, and bias configuration, then take a user-specified amount of frames to median-combine into a master dark frame, master flat field, and calculate a bad pixel map. The bad pixel locator algorithm uses bias/dark frames to create a histogram of all pixel intensities, fits a gaussian to the histogram, and marks pixels exceeding some σ threshold specified by the user as “bad,” creating a binary image where the offending pixels are 1s and all others are 0s. These calibrations can be bypassed by specifying the generation of default frames during the program’s execution; this creates dark, flats, and bad pixel frames of all zeros, ones, and zeros respectively.

During program execution, every frame is cleaned according to the two steps:

$$\text{Raw} = \text{BadpixelFilter}[\text{Raw}] \quad (33)$$

$$\text{Clean} = \frac{\text{Raw} - \text{Dark}}{\text{Flat}} \quad (34)$$

where the bad pixel removal replaces the locations in the frame marked as bad pixels with a 5x5 median filtered image of the same frame.

3.1.2 Deformable mirror - image plane registration

The program *dm_registration* uses the following method to determine the image registration

1. Take a background frame (optional)
2. Use the deformable mirror to create speckles at $\langle 0, \pm k\lambda/D \rangle$ and $\langle \pm k\lambda/D, 0 \rangle$, where k is determined by the user
3. Locate the artificial speckles on the (optionally background subtracted) image, either automatically or by having the user click on them
4. Centroid each spot using a 2-dimensional Gaussian surface
5. Use the centroid vectors to determine the image center, deformable mirror rotation, and λ/D
6. Remove the artificial speckles

To determine the image center, angle, and λ/D , the following formulae are used:

$$\lambda/D = \frac{\frac{1}{2} (|\vec{x}_3 - \vec{x}_1| + |\vec{x}_4 - \vec{x}_2|)}{2k} \quad (35)$$

$$\vec{x}_c = \frac{1}{4} \sum_{j=1}^4 \vec{x}_j \quad (36)$$

$$\theta = \frac{1}{4} \sum_{j=1}^4 \tan^{-1} [\vec{x}_j - \vec{x}_c] + c_j \quad (37)$$

where each vector \vec{x}_i refers to the centroid vectors in pixel coordinates, and the speckles are indexed by j , proceeding clockwise. The first equation simply expresses the fact that the speckles are expected to be at $k\lambda/D$, where k some number from 1 to N , $N\lambda/2D$ being the maximum controllable radial of the deformable mirror in the image plane. The second formula calculates the center of the control region as the element-wise mean of the spot centroids. The last formula calculates the rotation angle as the mean of each rotation angle calculated from each spot, with a correction $c = [0, \pi, 0, \pi]$ applied to account for the fact that each spot is 90 degrees away from the next when viewed from the center.

3.1.3 Deformable mirror intensity calibration

Another important calibration that needs to be done is the intensity calibration of the deformable mirror. Excluding chromatic elongation, each speckle has a characteristic size of λ/D .

1. Take a background; apply this to all further images take in the following steps
2. for k in $[k_{\min}, k_{\text{step}}, k_{\max}]$:
 - (a) Generate a speckle at $\langle \pm k\lambda/D, 0 \rangle$ and $\langle 0, \pm k\lambda/D \rangle$ at a fixed deformable mirror amplitude A
 - (b) Perform aperture photometry on each speckle using a circle of diameter λ/D
 - (c) Average the four photometric readings $\rightarrow I(k)$
 - (d) Remove speckles
3. Fit a simple and invertible function to $[k, I(k)]$, as listed below

where the functions we use to convert between intensity in the image plane and deformable mirror amplitude are

$$I[A, k] = \frac{A^2}{ak^2 + bk + c} \quad (38)$$

$$A[I, k] = \frac{I}{\sqrt{ak^2 + bk + c}} \quad (39)$$

where I is in units of counts measured from the aperture photometry of the speckle, and A is the amplitude in nanometers of the sinusoidal ripple on the deformable mirror. The use of a quadratic function was found to be a good choice, and the amplitude A influencing the intensity quadratically is motivated by the fact that the electric field in the pupil plane is related to electric field in the focal plane by Equations 5 and 16. The modulus-square of Equation 16 therefore is proportional to the intensity in the focal plane; showing that the amplitude of a ripple in the pupil plane is quadratically related to the intensity in the focal plane.

The choice of the amplitude A to use to perform the calibration, set by the user in a configuration file, is important to ensure an accurate result. A should be large enough interference between the artificial speckles and the underlying speckle field (which is always present at some level) does not bias the measurements, but small enough that the mirror response is linear. Generally, we found that about 5-10 times the speckle amplitude gave a reasonable performance, with data deviating from the fit at the 10% level. Exact calibration is not required, as the control loop (following section) typically applies a correction with a gain smaller than unity.

3.1.4 Control region definition

The control region refers to the area in the focal plane where speckles are suppressed. The control region may be an arbitrary shape, but shapes that subtend more than π radians when viewed from the center of the image (ie, “two-sided shapes”) will not be able to achieve as deep nulls as “single-sided” control regions, where both phase and amplitude errors can be removed. In this code, control regions are defined as binary images consisting of 1s where speckle suppression is required and 0s elsewhere.

We implemented both a general and automatic control region definition routine. The automatic program creates an annular control region between user-specified radials $k_{\text{in}}\lambda/D$ and $k_{\text{out}}\lambda/D$, for use in two-sided nulls. The general control region program takes an image of the speckle field and has the user select points. It then generates a polygonal control region with the points as vertices, using the ray-casting algorithm to determine which points are interior to the vertices and which are exterior. The vertices are saved in the configuration file for reasons that will be made clear later.

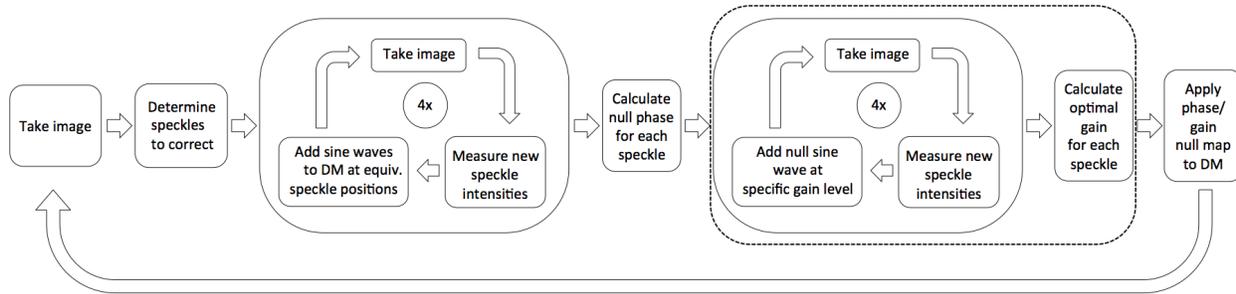


Figure 2: A graphical representation of the main control loop. The steps in the dashed section are optional. Image registration, control region definition, intensity calibration, etc are not shown here, and must be performed before this portion of the code is run for the first time.

3.2 Controls

The general operation of the speckle nulling code is presented below, with a graphical flow diagram shown in Figure 2.

1. Identify bright pixels in the image using a local maximum filter of a user-set size
2. Filter those points which are closer than a user-set threshold, keeping the brighter points
3. For $\phi = [0, \pi/2, \pi, 3\pi/2]$:
 - (a) For each remaining point, create a deformable mirror map corresponding to a speckle at its position and intensity, at phase ϕ
 - (b) Add the deformable mirror maps together to create a master phase map
 - (c) Apply the phase map to the deformable mirror
 - (d) Measure the intensities of the bright points using aperture photometry, as before
4. For each speckle, calculate the null phase using the intensities measured
5. Create a null phase map from the sum of all the speckle null maps
6. Either apply the null phase map to the deformable mirror with gain $g < 1$ or
 - (a) For $g' = [0, 0.33, 0.66, 1.1]$:
 - i. Apply the null phase map to the deformable mirror with gain g'
 - ii. Measure the intensities of the speckles
 - (b) Fit a parabola to each speckle's gain value vs. intensity; find the minimum of the parabola, call this the "null gain"
 - (c) Construct a "super nullmap" from individual deformable mirror maps using each null phases applied with each null gain
 - (d) Apply the super nullmap to the deformable mirror
7. Update the star centroid position (optional)

The first two steps deal with identification and filtering of speckles. A local maximum filter (just like a local median filter, except takes the maximum point) is used to identify bright points in the image. While this is not a particularly clever algorithm, it was found to be quite robust and sufficient for these purposes, compared with more sophisticated algorithms we tried using connected-components analysis that were extremely sensitive

to flux levels, speckle shapes, and required tuning of multiple parameters to achieve good performance. The second step can be viewed as removing points from consideration that correspond to spatial frequencies too close to simultaneously correct on the deformable mirror. This is very similar to filtering in spatial frequency space, though done in the image plane. Step (3) is the classic four-step phase shifting interferometry algorithm, where the “intensity” of the perturbing beam is matched to each speckle using aperture photometry. It should be pointed out that an exact match in flux does not have to be made, but having the artificial and natural speckle have the same flux increases the visibility, and hence phase accuracy. Step (4) calculates the phase using the regular four-step arctangent algorithm, and Step (5) simply adds the phase nullmaps for each speckle to create a null map for the entire region. The iteration may be terminated after uploading the nullmap to the deformable mirror (with some gain term to prevent overcorrection), but an alternative way to get somewhat better performance at the cost of four more images is to vary the gain of each speckle nullmap, measure the intensities, and then measure the optimal gain for each speckle. This prevents amplification of speckles when the phases are improperly measured. Additionally, it gives a good “stop” condition for the algorithm: when the optimal gains drop to zero, there is little point to running more iterations; the algorithm has converged.

The last step is also optional. For instruments that experience a changing gravity vector, it is often found that the image will drift on the detector over time. This can be a problem, as it will throw off the spatial-frequency/detector position calibration. As such, as the last step can re-solve for the center of the image using artificial speckles generated at the outer control radius of the deformable mirror, and update the control region and configuration file accordingly.

4. PERFORMANCE ON SELECTED INSTRUMENTS

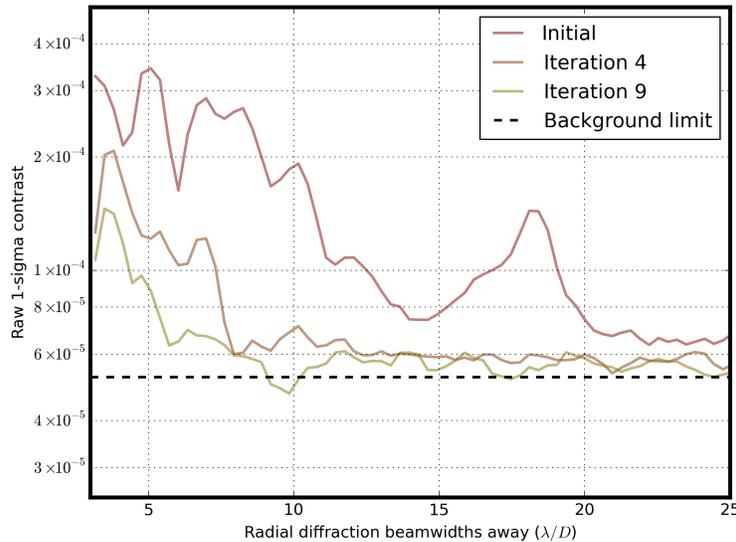
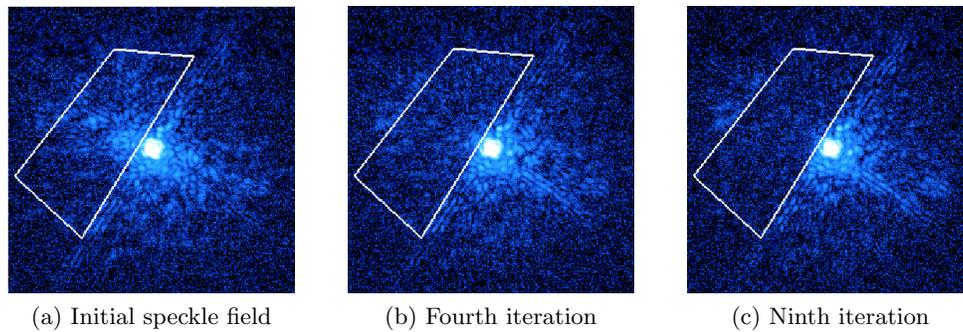
4.1 PHARO and P3K

PHARO⁸ is a near-infrared imager for the Palomar 200” Hale telescope. It operates at wavelengths from J to K band, including narrow-band filters like Fe II and Br- γ , and three grisms that deliver resolutions of about 1500 in J , H , and K bands. The detector is a Hawaii-I, with 25 μm pixels giving a plate scale of 25 mas/pixel or 40 mas/pixel, depending on the off-axis paraboloid used to bring the light to focus. PHARO has coronagraphic capabilities, with vector vortex coronagraphs installed for both H and K bands.

The P3K adaptive optics system⁹ is a high performance 3600 actuator adaptive optics for the Hale telescope. It works with many backend instruments, including PHARO. One advantage of P3K is that it has an image sharpening routine already in place to correct non-common path errors, using the Modified Gerchberg-Saxton (MGS) algorithm.¹⁰ With this procedure, images are taken before and after focus (by moving the input stimulus focus mirror), and these pairs of images can be used to determine the phase aberrations. The MGS algorithm takes a few iterations to converge, but can typically improve the internal Strehl ratios from the mid 70% level to better than 95%. Unfortunately, MGS tends to be insensitive to high-frequency speckle aberrations, whose signature gets washed out in the defocused images.

We implemented two forms of speckle nulling control for the P3K system; open and closed loop. For the open loop correction, the actuator voltages are used to change the shape of the mirror, then the wavefront sensor of the adaptive optics system reads out the centroids after the optimal shape of the mirror has been found. The system uses these new centroid offsets as the new setpoint when moving to the sky. In the closed loop mode, the centroid offsets are modified directly. (The conversion between centroid offsets and mirror shape must be determined through a separate procedure as part of the setup of the adaptive optics system. This matrix is stored in the AO system controls computer; we do not generate it ourselves.) In principle, the closed loop mode could be used to correct actively using a star as the stimulus, as is done with SCEXAO^{6,11} but the limitation at this point is the slow readout of PHARO, which takes a few seconds per frame.

In the examples presented below, rather than using the internal coronagraph of PHARO, we actually had a separate coronagraph installed in front of the instrument, delivering a better contrast and inner working angle. This instrument, the “Stellar Double Coronagraph,” or SDC, improves the level of residual diffraction in the image plane, as it corrects for the contrast-degrading effects of the secondary mirror that cannot be dealt with properly when using a single coronagraph and downstream pupil stop.¹² A full description of the Stellar Double Coronagraph may be found in Bottom et al. 2016.¹³ Apart from the improved starting contrast, the use of SDC



(d) Contrast in region

Figure 3: Coronagraphic speckle nulling in the dual-vortex mode using the internal white light source of the adaptive optics system. The value of λ/D is ~ 90 mas (a) The initial results of PSF correction using MGS still leaves many residual speckles in the focal plane. (b) Four iterations of speckle nulling remove most of the residual wavefront errors (c) Nine iterations get to within a factor of two of the detector read noise from 5 - 25 λ/D . The white polygon demarcates the control region, which is selected by mouse clicks in the half-region control mode. (d) Contrast improvement measured in the control region shows factors of 3-6 improvement, which are significant for companion detectability. The contrast curve is defined in the usual way, with the standard deviation (ie, 1σ) of surface brightness at each radial separation being used to generate the curve, and normalized by dividing by the peak flux of the non-coronagraphic PSF (not shown). The background limit is determined by the contrast in a region of the detector 100's of λ/D away. The preprocessing steps performed on the data only consist of dark subtraction and flat-fielding.

(or any other intermediate instrument) does not affect the speckle nulling code at all, as it simply modifies the adaptive optics system based on the final intensity information in the image plane, regardless of how it got there.

4.2 TMAS and P3K

The Ten-milliarcsecond Imager (Dekany et al., in prep.), or TMAS, is a visible light imager for the Palomar 200" Hale telescope. It operates at the wavelengths of 380 to 980 nm, and uses the NEO scientific CMOS detector from Andor technologies, with a 2.5k x 2.1k detector size and a readout time of up to 100 Hz. A number of filters are available on TMAS, including including broadband u, b, v, r, i and narrowband filters at 820 nm. TMAS operates behind the P3K adaptive optics system, providing diffraction-limited imaging at visible and

near-infrared wavelengths. Science applications of TMAS include time-resolved imaging of solar system objects, brown dwarfs, and stellar multiplicity studies.

TMAS uses off-the-shelf lenses and prisms to bring the corrected light from the adaptive optics system to focus on the detector, and as such, it has a high susceptibility to chromatic phase and amplitude speckle aberrations, especially when considering the short wavelengths of operation. Furthermore, TMAS currently does not possess an internal coronagraph (though there is the capability of installing one), so speckles are amplified by unsuppressed diffraction. Furthermore, the Modified Gerchberg-Saxton algorithm for static phase correction on the point-spread function core had not been implemented, so manual hand-tuning of Zernike polynomials on the deformable mirror was used to sharpen the point-spread function, which is less effective. All this causes the starting contrast to be much worse than in the case of PHARO or NIRC2, discussed in the next section.

The speckle nulling code was used in conjunction with TMAS in June 2015. Given the different camera, it was necessary to write a separate camera module compared to the one used with PHARO, but the adaptive optics module was used unchanged. It was found that the broadband filters were too wide and gave marginal results due to speckle chromaticity, so the narrowband filters were used. Results are presented in Figure 4. While the speckle field was decreased, particularly towards the outer working angles, the contrast improvement was modest.

4.3 NIRC2 and Keck AO

NIRC2[†] is the near-infrared imager on the Keck II telescope, designed to work with the facility adaptive optics system. The instrument uses an Aladdin-3 InSb array with imaging capabilities from 1 to 5 microns, with 27 μm pixels and pixel scales of 10, 20, and 40 mas/pixel. There are a number of different coronagraphic modes available on this instrument, which will be discussed below. The majority of directly imaged exoplanets have been discovered using NIRC2.

NIRC2 operates behind the Keck II adaptive optics system,¹⁴ which uses a 349 hexagonal actuator deformable mirror matched to the shape of the primary mirror. The adaptive optics system may be accessed through a Python wrapper[‡] to the KTL¹⁵ API, which permits on-the-fly modifications to the deformable mirror actuator positions, as well as control of multiple instruments, including NIRC2.

Before observing, it is customary to perform image sharpening routines to optimize the point-spread function. The image sharpening routines do not fix all the aberrations in the beam, but can give a high quality PSF with Strehl ratios over 90% in the L-band. Speckle nulling is then used after these calibrations have been performed.

While the speckle nulling code was first deployed using the conventional Lyot coronagraphs at *H*-band, and works well with that configuration, it is most often used with the new high-performance angular groove vortex coronagraph that was recently installed in NIRC2 (see B. Femenia-Castella et al., these proceedings). This *L*-band coronagraph gives a high level of diffraction control at close inner working angles, and benefits from the higher Strehl delivered at the longer wavelengths. For bright stars, the speckle noise dominates the error budget at these wavelengths.

The current implementation of the speckle nulling code on the Keck telescopes is still in progress, with only open-loop control being available—we transfer the mirror shape to the centroid offsets when switching from calibration fiber to star by averaging the wavefront sensor once the best mirror shape has been found. This is not so bad, as the typical readout time of NIRC2 for decent signal-to-noise on the speckles is on the order of 10 seconds; meaning that fast correction on-sky directly manipulating the centroid offsets of the adaptive optics system is some time away. Despite this, performance is very good, given the high quality of the adaptive optics system performance and of the optical path from the telescope to the detector.

We note that there is currently an additional step in open-loop mode due to the particular hardware of the Keck system. Here, the input source light is not carried through a single fiber, but two; one corresponding to visible and shorter wave infrared (and the wavefront sensor), and the other used to carry longer wave infrared, visible on the detector and used for speckle nulling. Both fibers are mounted on the same mechanical stage,

[†]<http://www2.keck.hawaii.edu/inst/nirc2/>

[‡]<http://spg.ucolick.org/KTLPython/>

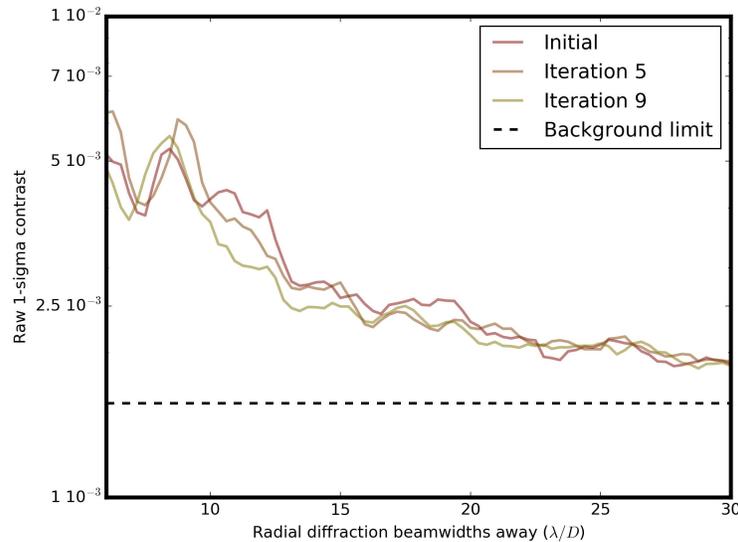
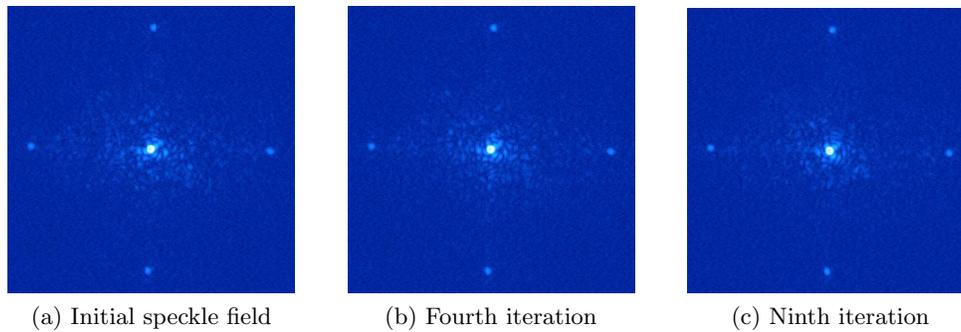
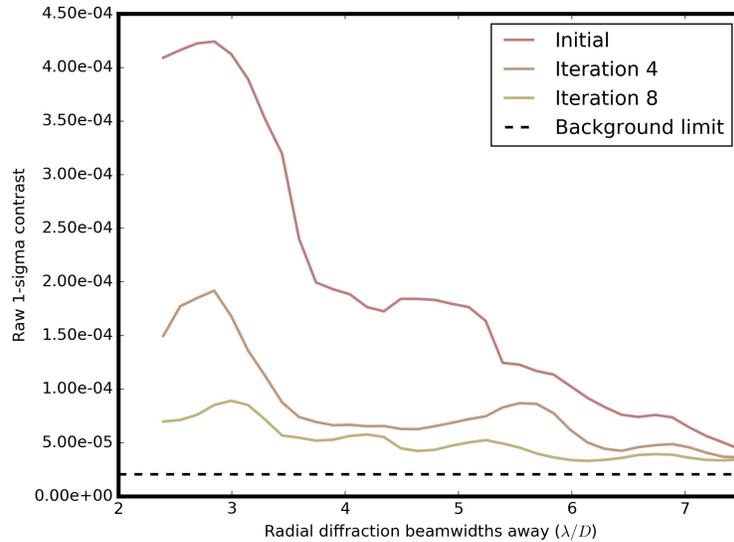
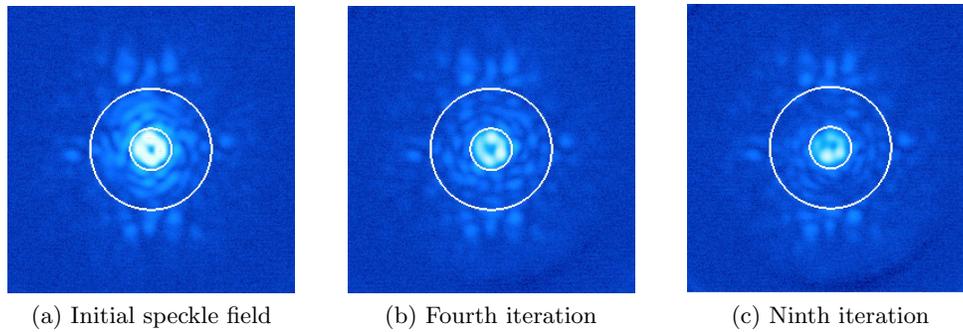


Figure 4: Non-coronagraphic speckle nulling with the instrument TMAS using the internal white light source of the adaptive optics system. The value of λ/D is ~ 30 mas. (a) The initial results of PSF correction using Zernike tuning leaves many speckle aberrations in the focal plane (b) Four iterations of speckle nulling remove (c) Nine iterations of speckle nulling. The contrast improvement in this case is a modest 20-40%, but could possibly have been increased by running for more iterations (at the time this experiment was performed, the control software did not provide automatic feedback about the contrast improvement, so we were not aware that there was much more improvement to be had). Some of the limitations here are due to the starting level of aberrations being far higher than in the longer-wavelength instruments, as described in the text. It is encouraging to note that more of the diffraction pattern is visible in the last image than in the first. The contrast curve is defined in the usual way, with the standard deviation (ie, 1σ) of surface brightness at each radial separation being used to generate the curve, and normalized by dividing by the peak flux of the non-coronagraphic PSF. The background limit is determined by the contrast in a region of the detector 100's of λ/D away. The preprocessing steps performed on the data only consist of dark subtraction and flat-fielding.

so when the code has converged to a good solution, we move the visible fiber to the precise position of the longer-wave infrared one before reading out the wavefront sensor. This step is meant to guarantee the same parts of the optics are illuminated by the wavefront sensor fiber and the speckle nulling fiber.

The contrast improvement using the L-band vortex were the best of all, especially when considering that a full annular region is corrected. This is very encouraging, as NIRC2 with the vortex coronagraph on Keck II is the most powerful combination of instruments in terms of planet detection sensitivity among those presented in this paper.



(d) Contrast in region

Figure 5: Coronagraphic speckle nulling in L-band on NIRC2, using the internal white light source of the adaptive optics system. The value of λ/D is ~ 80 mas. (a) The initial results of PSF correction using image sharpening leaves many speckle aberrations in the focal plane (b) Four iterations of speckle nulling remove (c) Nine iterations of speckle nulling. The white lines demarcate the control region, with the outer line at approximately the limit of the control bandwidth of the deformable mirror. The contrast improvement in this case is a factor of 2-3, and gets within a factor of 2-3 of the background limit. The contrast curve is defined in the usual way, with the standard deviation (ie, 1σ) of surface brightness at each radial separation being used to generate the curve, and normalized by dividing by the peak flux of the non-coronagraphic PSF. The background limit is determined by the contrast in a region of the detector 100's of λ/D away. The preprocessing steps performed on the data only consist of dark subtraction and flat-fielding.

5. DISCUSSION AND CONCLUSION

We have presented the design and performance of a speckle nulling code currently being used to remove non-common path aberrations in instruments on Palomar and Keck. The code is self-calibrating and written in a way that allows it to be ported to different instruments fairly quickly, with the main work being the need to write modules to read out the cameras remotely and change the shape of the deformable mirror of the adaptive optics system.

The code is not currently optimized for speed, as the instruments it is being used with have rather slow readout times, and premature optimization is the root of all evil. Informal tests of the speed indicate that it takes about 0.5 seconds per iteration, not counting camera and adaptive optics system I/O, with the vast majority of this being due to the display of the images and contrast curves on the screen for the benefit of the

user. It would seem to be possible to get at least another factor of ten or so in speed, which will be an urgent task in the near future, for new implementations on fast cameras.

A further planned implementation is on the DARKNESS instrument on Palomar, a photon-counting camera using microwave kinetic inductance detectors, with the capability of an effective readout rate of many kHz. To take advantage of the high speed and low noise of this instrument, it will be necessary to speed up the code and the P3K adaptive optics I/O to a level where the limitation is speckle photon noise. There are two steps to doing this which we are currently pursuing. First, an alternate channel is being built to access the adaptive optics system centroid offset vector directly, rather than through the normal P3K handshake procedure, which currently imposes a limit of 1 Hz. The new channel will be able to be accessed at a few hundred Hz, which is nearly as fast as the atmospheric speckles. This work is expected to be completed in the next few months, at which point code speedups will be necessary to keep up with DARKNESS. These can be done from both a hardware and software perspective. For example, the current speckle nulling control computer is over 10 years old and rather sad, and will likely impose a limit when trying to do the large matrix multiplications necessary for computing centroid offsets many times per second.

On the Keck side, we have a number of planned experiments. We intend to upgrade the fiber input system where a single photonic crystal fiber can deliver light from the visible to L band, removing the annoyance and possible introduction of non common path errors when co-aligning the visible fiber to the infrared one currently installed. Second, cosmetic and user-interface improvements will be implemented to make running the code more accessible. Finally, we intend to experiment with running speckle nulling on-sky with a star, despite the long readout times of NIRC2, as this would be a reasonable way to create very deep single-sided dark holes at known companion locations.

REFERENCES

1. C. Marois, D. Lafrenière, R. Doyon, B. Macintosh, and D. Nadeau, “Angular Differential Imaging: A Powerful High-Contrast Imaging Technique,” *ApJ* **641**, pp. 556–564, Apr. 2006.
2. C. Marois, D. W. Phillion, and B. Macintosh, “Exoplanet detection with simultaneous spectral differential imaging: effects of out-of-pupil-plane optical aberrations,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Proc. SPIE **6269**, p. 62693M, June 2006.
3. A. Give’on, “The electric field conjugation: A unified formalism for wavefront correction algorithms,” in *Frontiers in Optics 2009/Laser Science XXV/Fall 2009 OSA Optics & Photonics Technical Digest, Frontiers in Optics 2009/Laser Science XXV/Fall 2009 OSA Optics & Photonics Technical Digest*, p. AOWA3, Optical Society of America, 2009.
4. S. J. Thomas, A. A. Give’on, D. Dillon, B. Macintosh, D. Gavel, and R. Soummer, “Laboratory test of application of electric field conjugation image-sharpening to ground-based adaptive optics,” in *Adaptive Optics Systems II*, Proc. SPIE **7736**, p. 77365L, July 2010.
5. D. Savransky, B. A. Macintosh, S. J. Thomas, L. A. Poyneer, D. W. Palmer, R. J. De Rosa, and M. Hartung, “Focal plane wavefront sensing and control for ground-based imaging,” in *Adaptive Optics Systems III*, Proc. SPIE **8447**, p. 84476S, July 2012.
6. F. Martinache, O. Guyon, N. Jovanovic, C. Clergeon, G. Singh, T. Kudo, T. Currie, C. Thalmann, M. McElwain, and M. Tamura, “On-Sky Speckle Nulling Demonstration at Small Angular Separation with SCEAO,” *PASP* **126**, pp. 565–572, June 2014.
7. W. A. Traub and B. R. Oppenheimer, *Direct Imaging of Exoplanets*, pp. 111–156. Dec. 2010.
8. T. L. Hayward, B. Brandl, B. Pirger, C. Blacken, G. E. Gull, J. Schoenwald, and J. R. Houck, “PHARO: A Near-Infrared Camera for the Palomar Adaptive Optics System,” *PASP* **113**, pp. 105–118, Jan. 2001.
9. R. Dekany, J. Roberts, R. Burruss, A. Bouchez, T. Truong, C. Baranec, S. Guiwits, D. Hale, J. Angione, T. Trinh, J. Zolkower, J. C. Shelton, D. Palmer, J. Henning, E. Croner, M. Troy, D. McKenna, J. Tesch, S. Hildebrandt, and J. Milburn, “PALM-3000: Exoplanet Adaptive Optics for the 5 m Hale Telescope,” *ApJ* **776**, p. 130, Oct. 2013.
10. 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,” in *Adaptive Optics Systems II*, Proc. SPIE **7736**, p. 77365X, July 2010.

11. N. Jovanovic, F. Martinache, O. Guyon, C. Clergeon, G. Singh, T. Kudo, V. Garrel, K. Newman, D. Doughty, J. Lozi, J. Males, Y. Minowa, Y. Hayano, N. Takato, J. Morino, J. Kuhn, E. Serabyn, B. Norris, P. Tuthill, G. Schworer, P. Stewart, L. Close, E. Huby, G. Perrin, S. Lacour, L. Gauchet, S. Vievard, N. Murakami, F. Oshiyama, N. Baba, T. Matsuo, J. Nishikawa, M. Tamura, O. Lai, F. Marchis, G. Duchene, T. Kotani, and J. Woillez, "The Subaru Coronagraphic Extreme Adaptive Optics System: Enabling High-Contrast Imaging on Solar-System Scales," *PASP* **127**, pp. 890–910, Sept. 2015.
12. D. Mawet, E. Serabyn, J. K. Wallace, and L. Pueyo, "Improved high-contrast imaging with on-axis telescopes using a multistage vortex coronagraph," *Optics Letters* **36**, p. 1506, Apr. 2011.
13. M. Bottom, J. C. Shelton, J. K. Wallace, R. K. Bartos, J. Kuhn, D. Mawet, B. Mennesson, R. Burruss, and E. Serabyn, "Stellar Double Coronagraph: A Multistage Coronagraphic Platform at Palomar Observatory," *PASP*, in press.
14. P. L. Wizinowich, D. Le Mignant, A. H. Bouchez, R. D. Campbell, J. C. Y. Chin, A. R. Contos, M. A. van Dam, S. K. Hartman, E. M. Johansson, R. E. Lafon, H. Lewis, P. J. Stomski, D. M. Summers, C. G. Brown, P. M. Danforth, C. E. Max, and D. M. Pennington, "The W. M. Keck Observatory Laser Guide Star Adaptive Optics System: Overview," *PASP* **118**, pp. 297–309, Feb. 2006.
15. W. F. Lupton and A. R. Conrad, "The Keck Task Library (KTL)," in *Astronomical Data Analysis Software and Systems II*, R. J. Hanisch, R. J. V. Brissenden, and J. Barnes, eds., *Astronomical Society of the Pacific Conference Series* **52**, p. 315, Jan. 1993.