



CHICAGO JOURNALS



High-Contrast Imaging and Wavefront Control with a PIAA Coronagraph: Laboratory System Validation

Author(s): Olivier Guyon, Eugene Pluzhnik, Frantz Martinache, Julien Totems, Shinichiro Tanaka, Taro Matsuo, Celia Blain, Ruslan Belikov

Reviewed work(s):

Source: *Publications of the Astronomical Society of the Pacific*, Vol. 122, No. 887 (January 2010), pp. 71-84

Published by: [The University of Chicago Press](#) on behalf of the [Astronomical Society of the Pacific](#)

Stable URL: <http://www.jstor.org/stable/10.1086/649802>

Accessed: 06/02/2012 10:28

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press and Astronomical Society of the Pacific are collaborating with JSTOR to digitize, preserve and extend access to *Publications of the Astronomical Society of the Pacific*.

<http://www.jstor.org>

High-Contrast Imaging and Wavefront Control with a PIAA Coronagraph: Laboratory System Validation

OLIVIER GUYON

National Astronomical Observatory of Japan, Subaru Telescope, Hilo, HI 96720 and Steward Observatory,
University of Arizona, Tucson, AZ 85721; guyon@naoj.org

EUGENE PLUZHNIK

NASA Ames Research Center, Mountain View, CA 94035

FRANTZ MARTINACHE, JULIEN TOTEMS, AND SHINICHIRO TANAKA

National Astronomical Observatory of Japan, Subaru Telescope,
Hilo, HI 96720

TARO MATSUO

NASA Jet Propulsion Laboratory, Pasadena, CA 91109

CELIA BLAIN

University of Victoria, Victoria, BC V8W 3P6, Canada

AND

RUSLAN BELIKOV

NASA Ames Research Center, Mountain View, CA 94035

Received 2009 April 17; accepted 2009 November 5; published 2009 December 11

ABSTRACT. The Phase-Induced Amplitude Apodization (PIAA) coronagraph is a high-performance coronagraph concept able to work at small angular separation with little loss in throughput. We present results obtained with a laboratory PIAA system including active wavefront control. The system has a 94.3% throughput (excluding coating losses) and operates in air with monochromatic light. Our testbed achieved a 2.27×10^{-7} raw contrast between $1.65\lambda/D$ (inner working angle of the coronagraph configuration tested) and $4.4\lambda/D$ (outer working angle). Through careful calibration, we were able to separate this residual light into a dynamic coherent component (turbulence, vibrations) at 4.5×10^{-8} contrast and a static incoherent component (ghosts and/or polarization mismatch) at 1.6×10^{-7} contrast. Pointing errors are controlled at the $10^{-3}\lambda/D$ level using a dedicated low-order wavefront sensor. While not sufficient for direct imaging of Earthlike planets from space, the 2.27×10^{-7} raw contrast achieved already exceeds requirements for a ground-based extreme adaptive optics system aimed at direct detection of more massive exoplanets. We show that over a 4 hr period, averaged wavefront errors have been controlled to the 3.5×10^{-9} contrast level. This result is particularly encouraging for ground-based extreme-AO systems relying on long-term stability and absence of static wavefront errors to recover planets much fainter than the fast boiling speckle halo.

1. INTRODUCTION

An imaging system aimed at detection or characterization (spectroscopy) of exoplanets must overcome the large contrast between the planet and its star. This is particularly challenging for Earthlike planets, where the contrast is $\approx 10^{-10}$ in the visible and the angular separation is $0.1''$ for a system at 10 pc. Many coronagraph concepts have recently been proposed to overcome this challenge (see review by Guyon et al. 2006). Among the approaches suggested, Phase-Induced Amplitude Apodization

(PIAA) coronagraphy is particularly attractive. In a PIAA coronagraph, aspheric optics (mirrors or lenses) apodize the telescope beam with no loss in throughput. A PIAA coronagraph combines high throughput, small inner working angle ($2\lambda/D$ for 10^{-10} contrast), low chromaticity (when mirrors are used), full 360° discovery space, and full $1\lambda/D$ angular resolution. Angular resolution (size of the planet's point-spread function [PSF] in the image) is a critical performance parameter. Exoplanet imaging sensitivity, even if speckles have been perfectly removed, is usually background limited due to sky or thermal

emission (near-IR and mid-IR imaging from the ground) or zodiacal and exozodiacal light (direct imaging of Earthlike planets in the visible from space). The PIAA concept, originally formulated by Guyon (2003), has since been studied in depth in several subsequent publications (Traub & Vanderbei 2003; Guyon 2005; Vanderbei & Traub 2005; Galicher et al. 2005; Martinache et al. 2006; Vanderbei 2006; Pluzhnik et al. 2006; Guyon et al. 2006; Belikov et al. 2006; Guyon et al. 2009; Lozi et al. 2009), which the reader can refer to for detailed technical information.

In the first laboratory demonstration of the PIAA concept (Galicher et al. 2005), lossless beam apodization was demonstrated, and the field aberrations introduced by the PIAA optics were confirmed experimentally. In this first prototype, the PIAA acrylic optics lacked surface accuracy required for high contrast imaging, and since this experiment did not include active wavefront control, the high contrast imaging potential of the technique could not be demonstrated. In the present article, we report on results obtained with a new system that includes reflective PIAA optics and wavefront control. Our prototype combines the main elements/subsystems envisioned for a successful PIAA imaging coronagraph instrument, with the exception of corrective optics required to remove the strong off-axis aberrations introduced by the PIAA optics. This last subsystem has been designed and built for another testbed, and its laboratory performance is reported in a separate article (Lozi et al. 2009).

The overall system architecture adopted for our test is presented and justified in § 2. The design of the main components of the coronagraphs (PIAA mirrors, masks) is also described in this section. Wavefront control and calibration are discussed in § 3. Laboratory results are presented in § 4.

2. LABORATORY SYSTEM ARCHITECTURE

2.1. Plate Scale in a PIAA System

Two optical conventions exist to define plate scale (physical distance in meters on the focal plane per unit of angle on the sky, in radians): one is to follow the trajectory of the central ray in the optical system and measure its displacement on the focal plane as the source is moved on the sky, and the other is to define the plate scale as equal to $(F/D_{fp}) \times D$, where D is the telescope diameter and F/D_{fp} is the beam F -ratio at the focal plane. While the two conventions give identical results in conventional optical systems, they yield different values at the focal plane after the PIAA optics. This effect, due to the pupil distortion introduced by the PIAA optics, is well documented in previous PIAA-related publications, and leads to some confusion when comparing plate scale values, as both definitions have previously been used. In this article, we chose to avoid using either definition; instead, we adopt a convention where focal plane scale is physically defined relative to sky angle: $1\lambda/D$ in the focal plane is defined by the physical distance by which the

PSF photocenter moves when the source position on the sky is offset by $1\lambda/D$.

2.2. Coronagraph Architecture

The coronagraph architecture adopted is a hybrid PIAA (Pluzhnik et al. 2006), where beam apodization is shared between the aspheric PIAA mirrors (described in § 2.4) and a post-apodizer (described in § 2.5). The PIAA mirrors perform most of the apodization, but leave a small amount of excess light at the edge of beam (left at 0.85% of the surface brightness at the center of the beam), which is then removed by the apodizer. Thanks to this hybrid approach, the PIAA mirrors are more easily manufacturable (less aspheric) and the apodizer tolerances are relaxed (the apodizer is not absorbing light in the bright parts of the beam). The hybrid design also solves the problem of propagation-induced chromaticity (Vanderbei 2006), which would otherwise limit contrast at $\approx 10^{-7}$ in a non-hybrid system working in a 20% wide band. While this second benefit was not relevant in our monochromatic experiment, it is key to enable high-contrast direct imaging of exoplanets from space.

The cost in throughput and angular resolution due to the apodizer are small since the apodizer only removes light in the fainter edges of the remapped beam.

A high-contrast image is formed after the apodizer, where starlight is blocked by the focal plane mask. Since the upstream PIAA optics + apodizer have apodized the beam with little loss in telescope angular resolution, the focal plane mask is small, with a radius ranging from approximately $1\lambda/D$ on the sky for a 10^{-6} contrast goal to approximately $2\lambda/D$ on the sky for a 10^{-10} contrast goal. The focal plane mask is also part of the low-order wavefront sensor (Guyon et al. 2009) briefly described in § 3.2 which uses starlight reflected by the focal plane mask for accurate sensing of pointing errors and defocus.

The optical layout of the laboratory experiment is shown in Figure 1. The light source is a single mode fiber fed by a HeNe laser ($\lambda = 632.58$ nm), mounted on an X, Y, Z stage for control of the input tip/tilt and focus. The PIAA system (mirrors PIAA M1 and PIAA M2) creates a converging apodized beam. PIAA M2 is chosen as the pupil plane for the system, and lens l1 creates a small image of the pupil plane onto the apodizer. Lens l2 reimages the pupil plane on the deformable mirror, which is located in a $\approx f/60$ converging beam. Lenses l3 and l4 form a focal plane for the focal plane mask. The reimaging lens l5 is used to create a pupil plane and a focal plane. A Lyot mask is located in the pupil plane, but can be remotely moved out to allow the science camera (which is nominally in the focal plane) to move forward and acquire a direct pupil plane image. The focal plane mask reflects some of the light to the coronagraphic low-order wavefront sensor (CLOWFS) camera. A detailed description of this device is given in Guyon et al. (2009).

Once the number of lenses/mirrors and their relative positions was decided, a set of equations was coded to link together

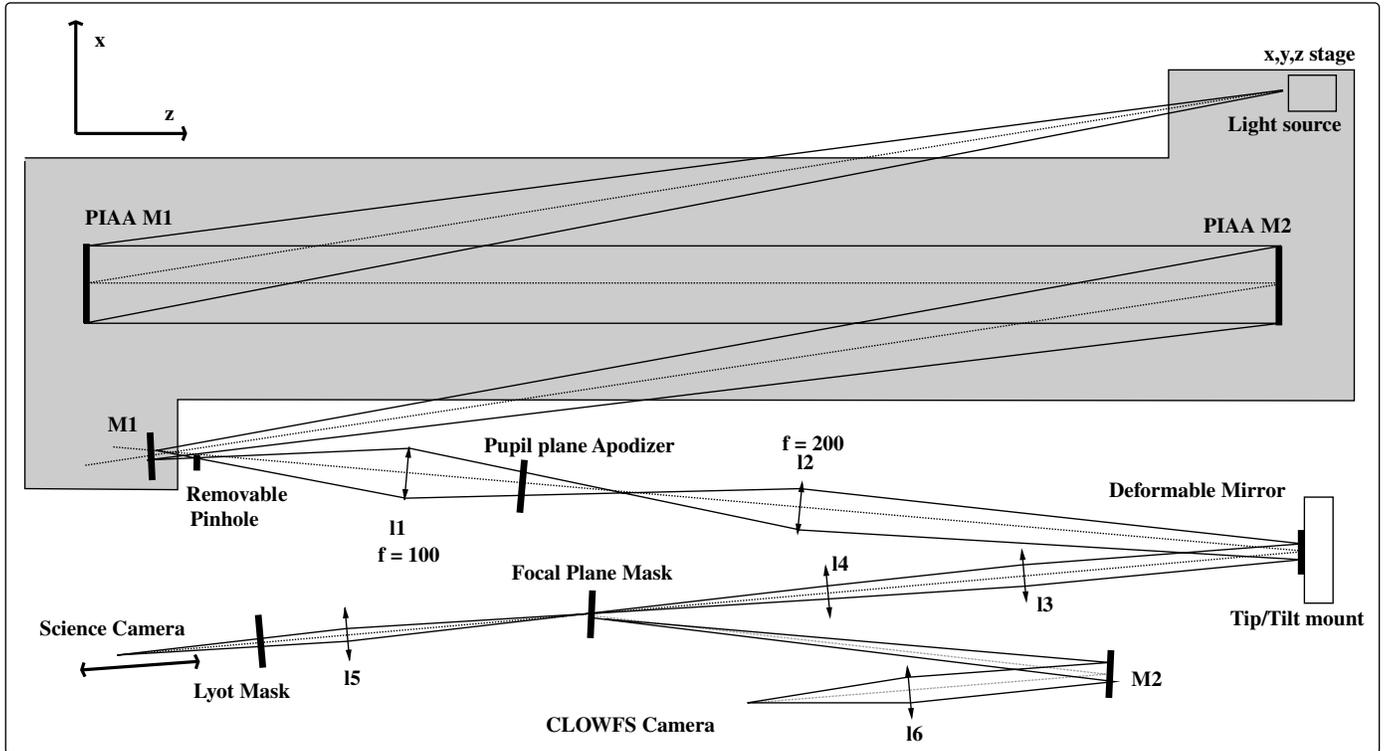


FIG. 1.—Optical layout of the laboratory PIAA coronagraph system. The *gray shaded* area shows the rigid PIAA bench on which the two PIAA mirrors are mounted. The light source is at the upper corner of the figure. The focal plane mask (near the *bottom, center*) separates light into the imaging channel and the coronagraphic low-order wavefront sensor (CLOWFS) channel.

the exact location of optical elements, their focal lengths (for lenses), the position of pupil and focal planes, the beam size on each optical element, and the plate scale in the focal planes. The equations contain all the hard constraints of the experiment (for example, the DM must be conjugated to the pupil plane apodizer). Each variable was given an allowed range (for example, beam size on the DM) or a set of allowed values (focal lengths of lenses constrained by what is available from vendors). A random search algorithm was then used to test many possible optical designs and identify the solutions that meet the criteria. This approach provided us with a flexible tool to explore design options. The same optimization code was also used to compute offsets in the position of several components during fine alignment of the system.

2.3. Wavefront Control Hardware and Architecture

A single 32×32 actuator MEMS deformable mirror (DM) is used for wavefront control and is located after the PIAA optics. The wavefront control subsystem is therefore fully decoupled from beam-shaping effects introduced by the PIAA optics. The PIAA optics simply deliver an apodized beam to the wavefront control subsystem, which operates independently of the technique used to apodize the beam (remapping vs. con-

ventional apodization). While this configuration is simpler than a configuration where the DM is ahead of the PIAA optics for wavefront control, we note that it does not offer as wide a field of view due to the magnification effect described in Guyon (2005).

Our laboratory demonstration was performed in monochromatic light. In this configuration, a single DM provides sufficient degrees of freedom to remove coherent diffracted light in one half of the field of view, regardless of how phase and amplitude aberrations in the beam are created. In a real coronagraphic instrument operating in broadband light, diffractive propagation between optics needs to be taken into account in design of the wavefront control architecture: manufacturing errors on optics introduce amplitude errors and wavefront chromaticity. In a high-contrast instrument, such errors can only be addressed with a multiple-DM configuration, where the DM locations are optimized to reduce residual wavefront errors over the spectral band used. Diffraction propagation effects between the aspheric PIAA optics surface would therefore need to be quantified in design of the wavefront control hardware.

Two challenging aspects of wavefront control in PIAA coronagraphs (understanding how beam shaping affects wavefront control, and polychromatic wavefront control) are therefore not addressed in our experiment.

2.4. Aspheric PIAA Mirror Design and Fabrication

The geometric remapping is performed by two highly aspheric mirrors, the PIAA mirrors. The role of the first mirror is mostly to project the desired amplitude profile on the second mirror, which is partially apodized with a faint plateau on the outside of the beam (see Fig. 2, *left*). The apodization profile is described in more detail in Pluzhnik et al. (2006), which includes a chromatic diffraction analysis of the PIAA optics used for our experiment. The PIAA M1 mirror acts as a converging element in the center (to concentrate more light in the center of the beam on PIAA M2) while light in the outside is diluted in a wide area of the beam on PIAA M2. This behavior explains the peculiar aspheric sag shown in Figure 2 on the right. The PIAA M2 mirror's role is to recollimate light to output a beam that is apodized but free of phase aberrations.

PIAA mirrors can be designed by solving a relatively simple differential equation when the input and output beams are collimated (the equation is given in Guyon 2003, and also in Traub & Vanderbei 2003 in a different form) or when the system is on-axis. In our laboratory experiment, the PIAA mirrors are focusing elements and the aspheric remapping shapes are added to off-axis parabolas. In this configuration, the PIAA mirror shapes cannot be derived from a simple differential equation, and they were designed by an iterative algorithm:

1. Initialization: The PIAA mirror shapes are computed by solving the differential equation for an on-axis system.
2. A constant slope is added to each of the PIAA mirrors. If there were no apodization, the mirrors obtained in Step 1 would be on-axis parabolas, and this slope would turn them into off-axis parabolas. In the PIAA system, however, adding this slope only leads to an approximation of the off-axis PIAA system.

3. A three-dimensional (3D) ray-tracing code is used to compute the beam phase and amplitude on the surface of a sphere centered on the output focus of the system immediately after reflection on PIAA M2.

4. The difference between the measured and desired beam amplitude on the sphere is used to update PIAA M1's shape by linear decomposition of this residual (using precomputed residuals obtained by adding Zernike polynomials on PIAA M1's shape).

5. The residual phase error measured on the sphere is compensated by changing PIAA M2's shape.

6. Return to Step 3 with the new mirror shapes.

This algorithm converges because changing PIAA M2's shape has little effect on the amplitude profile of the beam on the sphere, which is almost entirely a function of PIAA M1's shape.

The PIAA shapes were computed for a 75 mm beam diameter at the PIAA mirrors, a 1.125 m separation from the center of PIAA M1 to the center of PIAA M2, and a 190 mm offset between the PIAA M1 to PIAA M2 centerline and the input and output of the PIAA system (see Fig. 1). In the coordinate system shown in Figure 1, each mirror shape can be written as

$$z(x, y) = \text{OAP}(x, y) + \text{sag}(r) + \sum_i^j \alpha_i Z_i^j. \quad (1)$$

OAP(x, y) is the off-axis parabola which would be the PIAA mirror shape if no apodization was performed. It is a 1133 mm focal length OAP with a 190 mm off-axis distance from the center of the optical element. This shape is identical for the two PIAA mirrors although the orientation is different. $\text{sag}(r)$ is the apodization radial sag on each mirror. It is computed for

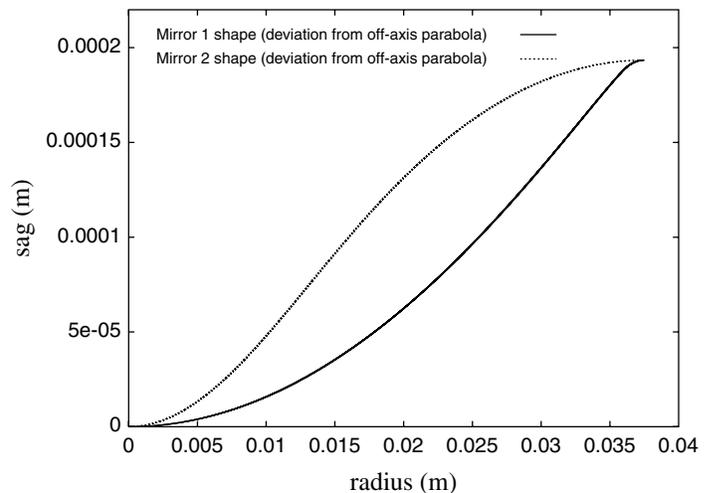
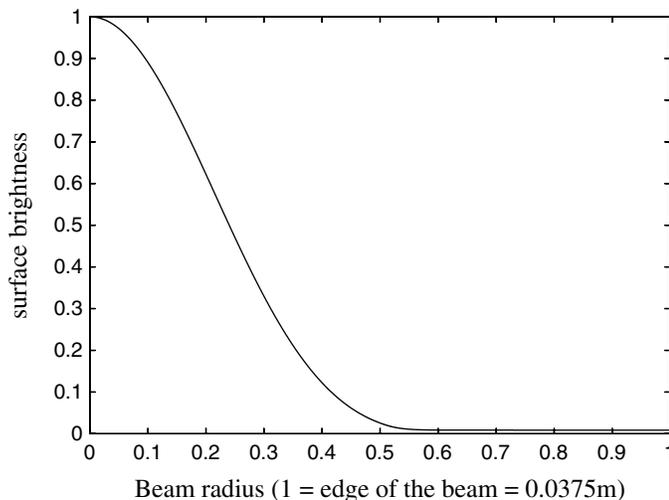


FIG. 2.—*Left*: Beam apodization profile used for the design of the PIAA system. In this hybrid design, the apodization is not complete, and the beam surface brightness at the edge of the beam is left at 0.85% of the center surface brightness. This extra light will need to be removed by a conventional apodizer. *Right*: Apodization radial sag term for PIAA M1 and PIAA M2 in an on-axis configuration. The narrow region at the edge of PIAA M1 has a strong localized curvature, and is the most challenging feature in manufacturing the PIAA system optics.

an on-axis system. A corrective term is added to account for the fact that the system is off-axis (tends to 0 for an on-axis system), decomposed as a sum of Zernikes polynomials up to radial order 7. For both PIAA mirrors, this correction is ≈ 95 nm rms (excluding tip/tilt).

The most challenging feature of the system is the small radius of curvature in the outer part of the beam on PIAA M1. While the hybrid design adopted mitigates this problem, the radius of curvature still reaches a minimum of 155 mm near the edge of the mirror, at 36.1 mm from the center of the mirror.

The PIAA mirrors were fabricated by Axsys Imaging Technologies. The mirror substrates were initially diamond turned according to the 3D prescriptions in the algorithm described, and then polished against computer-generated holograms (CGHs). PIAA M1 and PIAA M2 were then assembled on a rigid aluminum bench, aligned, and permanently fixed to the bench. The residual system wavefront error was then reduced to 0.04 waves rms by optical figuring of PIAA M2. Two sets of PIAA mirrors (4 mirrors total) were manufactured.

2.5. Postapodizer and System Throughput

In order to expedite manufacturing, the PIAA optics were designed to perform most, but not all, of the beam apodization required for high-contrast imaging. A more conventional apodizing scheme is therefore necessary to transform the beam profile at the output of the PIAA optics (solid curve in Fig. 3) into the desired beam profile (dashed curve in Fig. 3).

The postapodizer was designed in transmission, with a series of narrow opaque rings blocking light. The position and width of the rings is optimized to best approximate the ideal continuous apodization profile shown in Figure 3 as the curve labeled *apodizer transmission*. Several constraints were imposed on the design to ensure manufacturability: no ring should be less than $0.8 \mu\text{m}$ wide and the gap between consecutive opaque rings should be no less than $5 \mu\text{m}$. The resulting design is composed of 109 opaque rings for a total apodizer diameter of 3.815 mm (defined by the outer edge of the last opening between opaque rings). The apodizer was manufactured by lithography on a transmissive substrate.

The postapodizer throughput over the 3.815 mm diameter is 96.9 %, but due to the narrow rings in the apodizer, some of the light transmitted is diffracted at large angles. The effective throughput of the apodizer is 94.3%, and would be equal to the throughput if the apodizer were continuous instead of binary. The full system throughput can therefore reach 94.3% (excluding losses due to coating) provided that the telescope pupil size on PIAA M1 is adjusted to the apodizer diameter. In practice, the telescope pupil should however be made slightly larger to allow for pupil centering errors, and in very high contrast applications (space coronagraphy), to mitigate possible edge-ringing effects due to Fresnel propagation.

The apodizer throughput was measured by inserting a pinhole in the PIAA output focus and moving the science camera in the pupil plane. The $1 \mu\text{m}$ pinhole is used to deapodize the beam at the expense of a very low throughput. Figure 4 shows both the

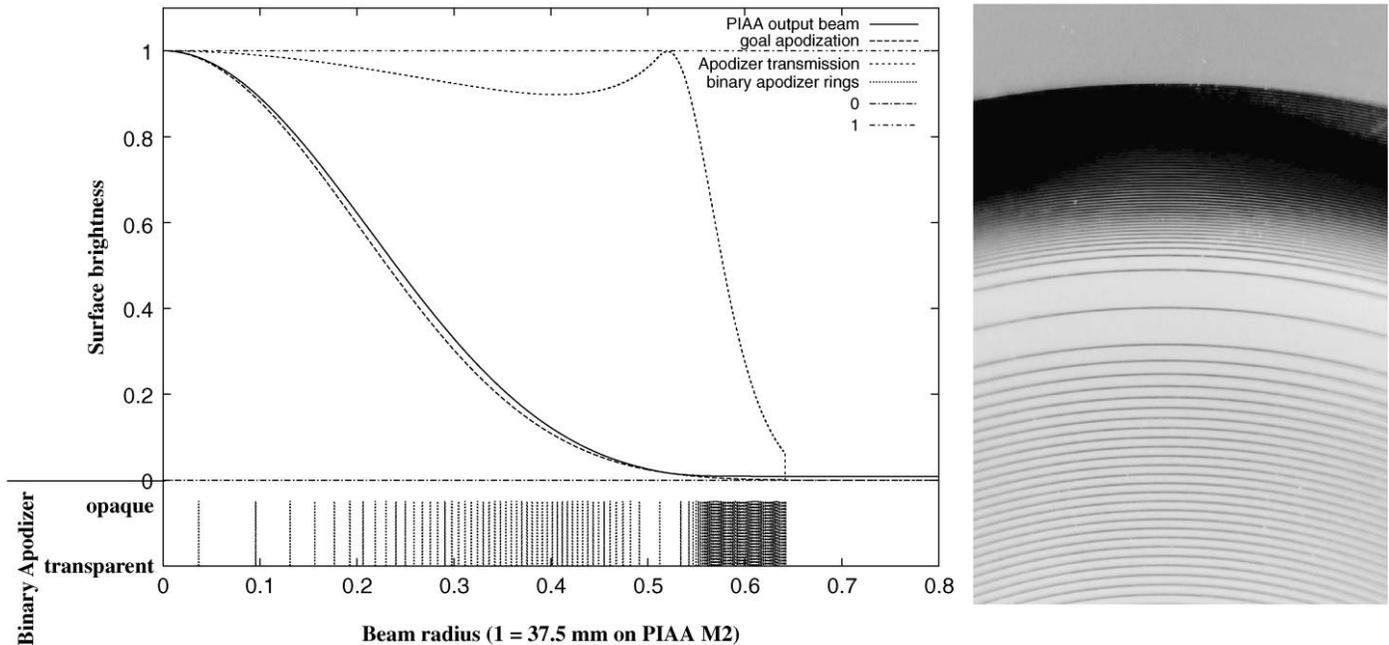


FIG. 3.—*Left*: Apodizer design. *Right*: Microscope image of the outer part of the apodizer (image covers approximately 1 mm vertically). The outer edge of the apodizer (last transmissive ring) is at 1.9 mm radius. The width of the individual opaque rings is $0.8 \mu\text{m}$.

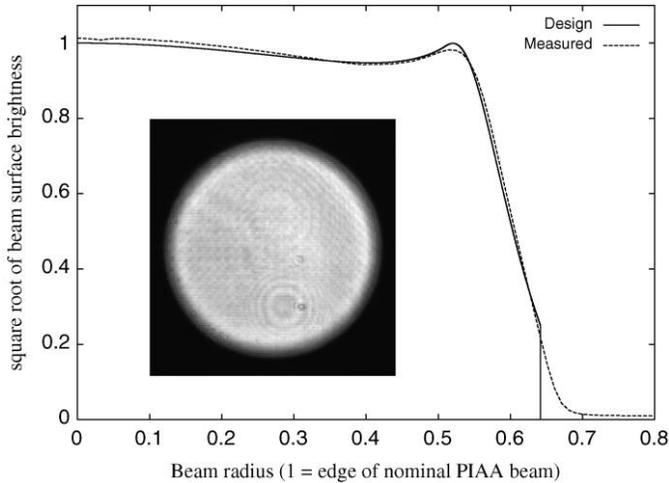


FIG. 4.—Laboratory measurement of the apodizer throughput. The designed and measured radial throughputs are compared.

measured apodizer profile and the designed apodization. The residual difference between the two curves is due to the finite size of the pinhole (the beam before the apodizer is slightly apodized, so the measured profile is slightly too bright in the center) and the finite angular resolution of the pupil reimaging (the sharp edge of the apodizer is blurred).

2.6. Focal Plane and Lyot Masks

The focal plane mask, shown in Figure 5, is used in transmission. The central part of the mask blocks the bright central PSF core. The radius of this nontransmissive central zone defines the inner working angle of the coronagraph, which is $1.65\lambda/D$ in our experiment. A clear zone transmits the science field to the science camera. The shape of this clear zone of the focal plane mask is chosen to exclude regions of the focal plane where the wavefront control system cannot remove diffracted light. Since our experiment uses a single deformable mirror, diffracted light can only be controlled over half of the field of view. The clear opening in the mask is therefore D -shaped with an outer radius imposed by the DM actuator sampling. A slightly larger rectangular zone could also have been adopted, but would have imposed the rotation angle of the focal plane mask. The reflective ring, extending from $0.8\lambda/D$ to $1.65\lambda/D$, sends some of the starlight to the low-order wavefront sensor (LOWFS) camera. We note that the opaque zones of the mask are not fully opaque due to manufacturing considerations: their transmission and reflection are, respectively, $\approx 10^{-5}$ and $\approx 10\%$.

A Lyot mask is located in the pupil plane between the focal plane mask and the science focal plane. This mask is designed to block all light outside the geometrical pupil and transmit all light within the pupil. It therefore has no effect on the nominal system throughput, and its role is to ensure that the scattered light reaching the focal plane camera does not contain light outside the

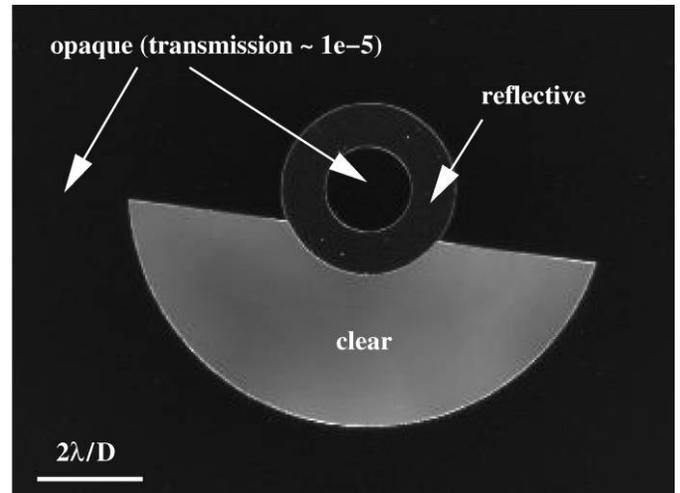


FIG. 5.—Microscope image of the focal plane mask.

pupil. Although correcting for such light is theoretically possible if it is coherent and within the control radius imposed by actuator sampling on the DM, it requires an accurate model of the coronagraph that can predict how light outside the pupil is affected by DM actuator positions. The Lyot mask was made by drilling a small hole in an aluminum plate, and its diameter is slightly smaller than the pupil size to account for alignment tolerances.

Both the focal and Lyot masks are on motorized stages and can be removed from the beam.

3. WAVEFRONT CONTROL

As described in § 2.3, wavefront control in our experiment is performed after the PIAA optics, which allows full decoupling between the pupil remapping introduced by the PIAA optics and the wavefront control algorithms. The wavefront calibration and wavefront control routines used in our experiment are therefore not specific to PIAA—they simply take as an input the apodized beam from the PIAA optics. We present these routines in this section as they are an essential part of the experiment, and the calibration techniques developed for this experiment are used to quantify the coronagraph’s performance beyond the raw contrast achieved. Section 3.1 describes how the wavefront is first flattened and the DM response is calibrated in a noncoronagraphic mode (no focal plane mask). Section 3.2 briefly describes the low-order wavefront control loop, which is detailed in a separate article. In § 3.3, the main wavefront control loop is described along with the calibrations used to measure the coronagraph performance beyond the raw contrast.

3.1. Initial Calibration Loop (without Focal Plane Mask)

Initial calibration is performed using conventional phase diversity with no focal plane mask: images are acquired with the science camera in six positions regularly spaced from the focal

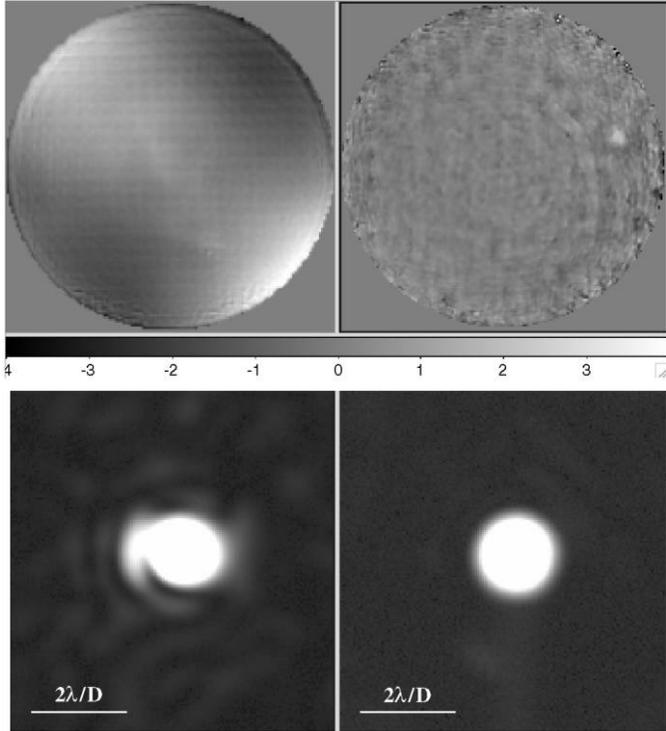


FIG. 6.—Pupil phase (*top*) and focal plane image (*bottom*) when the deformable mirror is powered off (*left*) and set to its nominal position after calibration (*right*). A malfunctioning actuator is visible on the right side of the beam.

plane to the pupil plane. An iterative Gerchberg-Saxon algorithm is used to reconstruct the pupil plane complex amplitude. As shown in Figure 6 (top left), the beam quality is initially quite poor, with a large amount of astigmatism. The corresponding focal plane image is shown in Figure 6, bottom left.

This phase diversity routine is repeated N times ($N \approx 10$), with a different set of DM voltages applied for each phase diversity measurement sequence. The N phase maps obtained and the N DM voltage maps used to obtain them are then used to constrain a model of the DM response consisting of seven parameters: Geometrical correspondence between the DM and the pupil image (4 parameters: x and y shift, scale, and rotation) and physical constants describing the DM behavior (3 parameters: width of the actuator influence function, DM displacement for 100 V applied, and power index α in the displacement to voltage relationship with displacement $\propto V^\alpha$). The result of the DM calibration can then be used to flatten the wavefront measured and produce a sharp focal plane image (Fig. 6, right).

This initial calibration is a necessary preliminary step for the high-contrast wavefront control, which needs (1) a knowledge of the starting point (typically less than 1 radian error on the wavefront) and (2) a good understanding of how DM commands affect the pupil plane phase. The quality of the DM calibration is shown in Figure 7 for a large DM offset (60 nm rms in this example). The difference between the measured and simu-

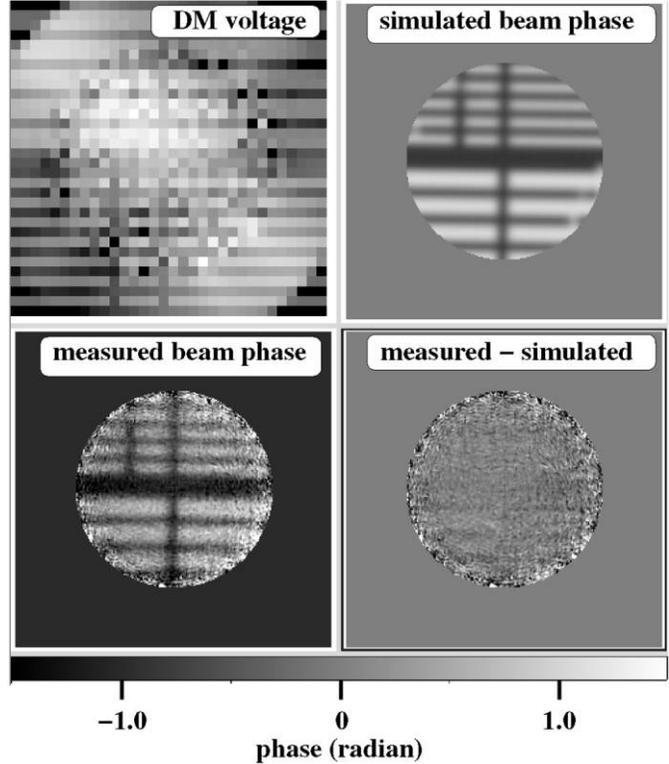


FIG. 7.—Result of the automatic DM calibration. The DM is driven to produce a recognizable pattern on top of the flat wavefront (DM voltages, *top left*). The DM model is used to convert the voltages into a simulated beam phase (*top right*). The difference between the measured beam phase (*bottom left*) and the simulated beam phase is shown in the *bottom right* panel. All images except the DM voltage map are shown with the same linear phase scale.

lated beam phase binned to the actuator size is 6 nm rms in the central 75% radius of the pupil (phase measurement in the outer part of the pupil is noisier due to lack of flux), corresponding to a 10% relative accuracy. The DM model relative accuracy is better for smaller displacements.

3.2. Low-Order Wavefront Errors

Low-order wavefront errors are measured by reflecting a portion of the bright starlight masked by the coronagraph focal plane into a dedicated camera. A detailed description of this low-order wavefront sensor (LOWFS) can be found in Guyon et al. (2009). The LOWFS signal is used to simultaneously drive the deformable tip/tilt and the source position ahead of the PIAA optics. A key feature of the LOWFS is the ability to separate pointing errors (pre-PIAA tip/tilt) from post-PIAA tip/tilt, which is essential to maintain high contrast: even a small pre-PIAA tip/tilt creates diffraction rings outside the IWA of the coronagraph, and pre-PIAA tip/tilt errors cannot be compensated for by post-PIAA tip/tilt.

When the low-order loop is closed, the measured residual pointing error is $10^{-3}\lambda/D$, and is therefore small enough to be negligible in the scattered light error budget shown in § 4.3.

A more detailed description of the design, calibration, control algorithm, and performance of the LOWFS in our experiment is given in Guyon et al. (2009).

3.3. High-Order Wavefront Control Loop

Coherent scattered light in the clear opening of the focal plane mask is measured by phase diversity introduced on the DM. A series of focal plane images, each acquired with a slightly different DM shape, is used to reconstruct the complex amplitude and coherence of the scattered light. The high-order wavefront control loop uses a linearized representation of the system in focal plane complex amplitude, as described in the electric field conjugation (EFC) approach proposed by Givéon et al. (2007). The wavefront control loop is shown in Figure 8, and is built around the EFC approach.

Prior to starting the loop, a model of the coronagraph is used to compute how each actuator motion affects the complex amplitude in the focal plane. Since this relationship is linear for small displacements, this model is stored as a complex amplitude system response matrix M_{resp} (shown on the righthand side of Fig. 8) of size n by m , where n is the number of DM actuators (ignoring actuators outside the pupil) and m is the number of pixels in the high-contrast region of the focal plane. M_{resp} is the linear operator that establishes the relationship between deformable mirror actuator displacements $\delta\text{DM}(u, v)$ (sampled in n points) and the corresponding complex amplitude change $\delta f(x, y)$ (sampled in m points) in the focal plane:

$$\delta f(x, y) = M_{\text{resp}} \delta\text{DM}(u, v). \quad (2)$$

M_{resp} is therefore computed by moving each actuator of the DM in the coronagraph model and storing the corresponding change in focal-plane complex amplitude in a column of M_{resp} .

3.3.1. Loop Initialization

For each iteration k of the loop, the first step in the wavefront sensing process is to acquire an image I_0^k with the DM shape set at the best known position for high-contrast imaging (upper left corner of Fig. 8). This first image is then used to choose the shapes to apply on the DM to optimize measurement accuracy and sensitivity.

3.3.2. Considerations for the Choice of Wavefront-Sensing DM Shapes

Residual light in the coronagraphic focal plane is measured by adding N known wavefront errors on the nominal deformable mirror shape. For each wavefront error added, a coronagraphic image is taken. Intensity variations between the N images encode the complex amplitude of the residual light that needs to be measured and removed. In this section, we discuss how to choose the N wavefront errors that are sent to the DM for sensing.

These N DM displacements are denoted $\delta\text{DM}_i^k(u, v)$, with $i = 1 \dots N$ (the index k denotes the wavefront control loop iteration). The complex amplitude added to the focal plane by each of the DM displacements is given by equation (2):

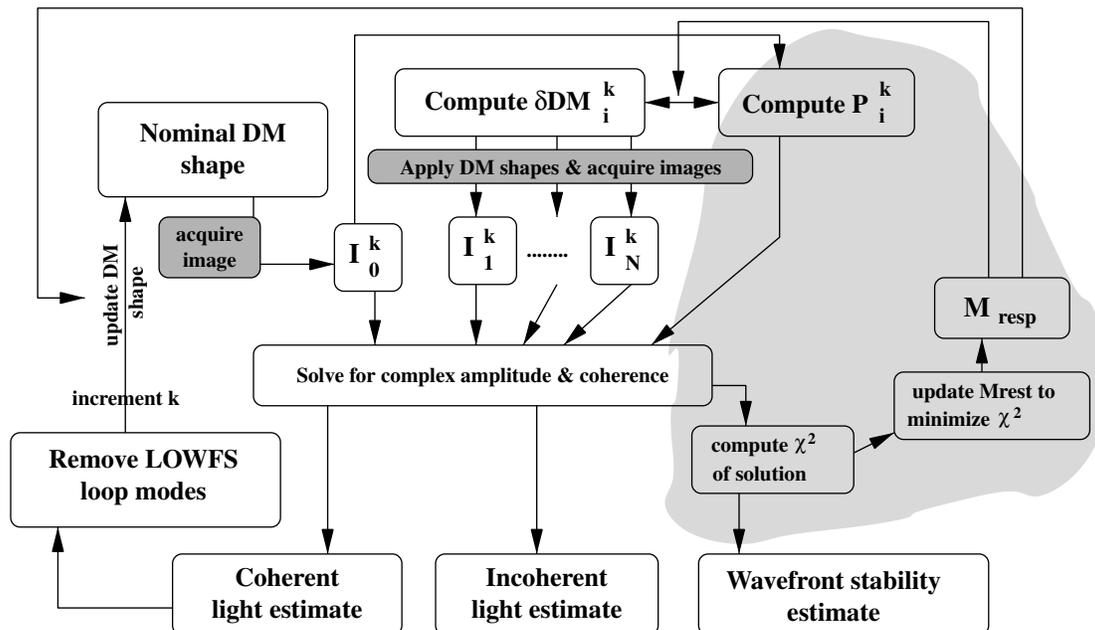


Fig. 8.—High-order wavefront control loop, showing both the main loop and the system response matrix optimization loop (light shaded area). The two dark shaded boxes indicate image acquisition, which, in the simulation mode, can be replaced with a simulated image acquisition using a model of the experiment and DM response.

$$P_i^k(x, y) = M_{\text{resp}} \delta \text{DM}_i^k(u, v). \quad (3)$$

These complex amplitude functions are referred to as wavefront-sensing probes. The amplitude of the probes must be carefully chosen: if the probes are too strong, the measurement is too sensitive to errors in the DM calibration; if they are too weak, the measurement is contaminated by photon noise, readout noise, and small variations in the incoherent scattered light. As a guideline, it is therefore best to choose these DM offsets so that the additional light (the complex amplitude focal plane probes) is approximately as bright as the light to be measured. Finally, randomly modulating the probes can mitigate the effect of calibration errors.

3.3.3. Wavefront-Sensing Probes

The first probe P_1^k is chosen to satisfy, for each pixel (x, y)

$$|P_1^k(x, y)|^2 = \alpha_0 + \alpha_1 I_0^k(x, y), \quad (4)$$

where $I_0^k(x, y)$ is the image acquired with DM shape DM_0^k . If $\alpha_1 = 0$, this constraint will force the DM shape to add a uniform coherent background of contrast α_0 in the focal plane, while if $\alpha_0 = 0$ and $\alpha_1 = 1$, it will drive the DM shape to add a speckle map with the same intensity as in the I_0^k image. The phase of this probe is not constrained, and is chosen by an iterative scheme to best satisfy equation (4), with δDM_1^k as the free parameter. We note that an exact solution to equation (4) may not exist within the DM space, but the following algorithm yields a good approximate solution:

1. Initialization: P_1^k is computed from equation (4) with a zero phase.
2. Singular value decomposition (SVD) of M_{resp} is used to estimate δDM_1^k from P_1^k . Eigenvalues close to zero are rejected in the SVD to improve stability.
3. δDM_1^k is clipped to avoid large DM displacements.
4. Equation (3) is used to recompute P_1^k from the clipped δDM_1^k .
5. The amplitude of P_1^k is updated to satisfy equation (4), but its phase is left unchanged from Step 4.
6. Return to Step 2 until the iterative algorithm is stopped.

This iterative algorithm simultaneously produces the probe P_1^k and the corresponding DM displacements.

The second probe is chosen so that, at each point (x, y) in the focal plane, its amplitude is identical to the first probe, but its phase is offset by $\pi/2$:

$$P_2^k(x, y) \approx iP_1^k(x, y). \quad (5)$$

This $\pi/2$ phase offset maximizes the WFS sensitivity if the dominant sources of noises are photon noise and readout noise (Guyon 2005). We note that if all DM actuators are functioning, there is a perfect solution to this equation, which can be ob-

tained by shifting each spatial frequency of the DM_1^k map by $\pi/2$. Images acquired with these first two probes, together with the image I_0^k , would be sufficient to solve for wavefront errors if light in the focal plane is fully coherent, but at least one more probe is needed to unambiguously measure light coherence, and more probes can also provide the redundancy required for implementation of the diagnostic tools described in §§ 3.3.4 and 3.3.5.

Two additional probes have been chosen to be $P_3^k = -P_1^k$ and $P_4^k = -P_2^k$, with exact solutions $\delta \text{DM}_3^k = -\delta \text{DM}_1^k$ and $\delta \text{DM}_4^k = -\delta \text{DM}_2^k$ respectively.

In our laboratory experiment, we chose to also add five additional probes with random uncorrelated DM shapes of similar amplitude as the DM displacements obtained for probes 1 to 4. The mean displacement amplitude per actuator is measured on probes 1 to 4, and random displacement maps are produced with the same amplitude. These additional probes are not required for wavefront reconstruction (probes 1 to 4 provide sufficient information), but, as described in the next sections, are added to allow for calibration of M_{resp} , measurement of incoherent light in the system, and measurement of coherent light variation during the sequence of N exposures.

3.3.4. Solving for Complex Amplitude and Coherence

For each pixel (x, y) of the focal plane detector, the complex amplitude $A(x, y)$ of the coherent light leak and the intensity $I(x, y)$ of the incoherent light leak is estimated by solving the following set of equations:

$$I_i^k(x, y) = |A(x, y) + P_i^k(x, y)|^2 + I(x, y) \quad (6)$$

for $i = 0 \dots N$, with $P_0^k(x, y) = 0$ (image acquired with nominal DM shape).

This set of equations has three unknowns and is therefore overconstrained for $N > 2$. With $N = 9$ adopted in our experiment, $A(x, y)$ and $I(x, y)$ are chosen to minimize:

$$\chi^2(x, y) = \sum_{i=1 \dots N} \{I_i^k[x, y] - [|A(x, y) + P_i^k(x, y)|^2 + I(x, y)]\}. \quad (7)$$

In noiseless ideal simulations, the residual χ^2 should be null. The residual χ^2 in the laboratory experiment includes errors due to:

1. Photon and readout noise in the $I_i^k(x, y)$ measurements
2. Variations in the light leaks during the measurement sequence. When solving for this set of equations, we assume that $A(x, y)$ and $I(x, y)$ are static, but if they vary, χ^2 will increase.
3. Systematic model errors (errors in M_{resp}), which lead to errors in the estimation of the values of $P_i^k(x, y)$. If M_{resp} is wrong, then the DM command sent will not produce the expected $P_i^k(x, y)$.

For convenience, we have scaled χ^2 in coronagraphic contrast units. This scaling is performed by measuring the uncorrelated noise that would need to be added to $A(x, y)$ between frames to reproduce the observed value of χ^2 . In this unit, the observed χ^2 is approximately 4.5×10^{-8} (see § 4).

The first contribution (photon + readout noise) has been computed to be a small part of the χ^2 observed. We observed that increasing α_1 (see equation (4)) above 1 has little effect on the residual χ^2 , also independently suggesting that χ^2 is not due to detector readout noise.

3.3.5. System Response Matrix Optimization

One key output of the wavefront control loop is the estimation of coherent light leaks (which should be used to compute the DM correction to apply for the next iteration), incoherent light leaks (which the DM can do nothing about), and the measurement of the wavefront stability during the measurement sequence. All these quantities depend upon a reliable estimation of M_{resp} . It is possible, for example, if M_{resp} is wrong, to obtain a low value of the residual coherent light and think that the system has converged to a good contrast value, while in fact a significant amount of coherent light remains. This last issue is mitigated, but not entirely addressed, by continuously varying the probes (as this error is a function of the probes chosen, and will average to zero in the linear approximation used in this work if the probes are randomly chosen). An error in M_{resp} would first appear as a large χ^2 value for the solution of equation (6).

To address this, we have added a M_{resp} optimization loop within our control loop. For each iteration k , the derivative of

$$\chi^2 = \sum_{x,y} \chi^2(x, y) \quad (8)$$

with each element of M_{resp} is computed (this is a total of $2 \times n \times m$ derivatives, as the derivative is computed for the real and imaginary parts of each element of the M_{resp} matrix). This derivative is computed from equation (7) by replacing P_i^k with $M_{\text{resp}} \delta \text{DM}_i^k$ (equation (3)).

At each iteration of the wavefront control loop, M_{resp} is then slightly modified in order to reduce χ^2 : for the 10% of M_{resp} elements showing the largest derivative against χ^2 , the value of the matrix element is moved in the direction indicated by the derivative by 0.1% of the rms value of elements in the corresponding column of M_{resp} . This algorithm was first tested on simulated data with an initial M_{resp} estimate that was different from the actual M_{resp} used in the simulation for computing the images. This test showed that M_{resp} did converge toward the true M_{resp} , and that the χ^2 value decreases as a result. Convergence is very slow, due to the large number of coefficients in the M_{resp} matrix, requiring several hundred iterations before a significant improvement in χ^2 is observed.

3.3.6. Correction Applied to the DM

The linear electric field conjugation (EFC) algorithm (Givéon et al. 2007) is used to cancel coherent scattered light. This algorithm uses the linearized coronagraph model which is also used for the measurement step described in § 3.3.4. The system response matrix is inverted to build a control matrix which is multiplied to the coherent light estimate to yield the DM shape offset to be applied. Regularization schemes proposed by Givéon et al. (2007) were used to improve the loop stability and convergence speed.

4. LABORATORY RESULTS

4.1. Alignment and Apodization Measurement in the Pupil Plane

The pupil plane apodization map is measured by placing the science camera in the pupil plane. In our experiment, the pupil is conjugated simultaneously to the PIAA M2 mirror, the apodizer, and the DM. Alignment is necessary to ensure that these three planes are conjugated and that their relative scales are correct. The camera positions for which conjugation to these planes is achieved are measured and the corresponding pupil scales are derived from the images. These six numbers are then fed to an optimization routine which computes the offsets to be applied to all movable optical elements after PIAA M2 to meet the conjugation and scale requirements. A few iterations of this sequence were sufficient to converge.

In the fine alignment step, the apodizer alone is moved. The pupil image is compared to a simulated pupil image where the relative scale and lateral offset between the PIAA apodization and the apodizer transmission map are free parameters. The values of these three parameters that give the smallest residual difference are then used to guide fine alignment of the apodizer. Fine tuning of the scale between the apodizer and the PIAA apodization is possible because the beam at the apodizer is non-collimated: apodizer motion along the optical axis changes this scale. Figure 9 shows both the pupil apodization profile measured after alignment and the theoretical profile.

4.2. Imaging with a Noncorrected PIAA System

Figure 10 shows the system on-axis PSF in imaging mode (no coronagraph focal plane mask). The on-axis PSF is similar to an Airy function without the Airy rings beyond $1.22\lambda/D$.

While the on-axis image is sharp and exhibits high contrast, our laboratory system did not include the inverse PIAA optics necessary to correct for the strong off-axis aberrations introduced by the forward PIAA optics. These inverse optics do not need to be of coronagraphic quality, and can be a small set of lenses. A laboratory demonstration of wide-field correction with inverse PIAA optics is described in a separate article (Lozi et al. 2009).

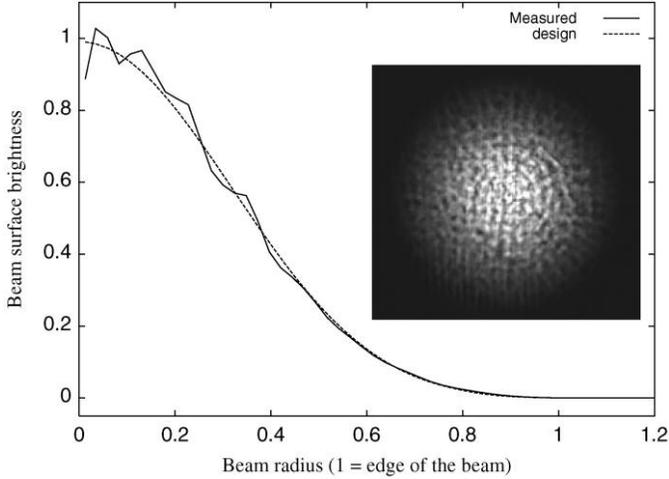


FIG. 9.—Measured apodization profile, compared the to the beam profile the experiment was designed to deliver. A 2D image of the beam is also shown.

Without inverse PIAA optics, the image of an off-axis source rapidly changes shape as the source moves away from the optical axis. Figure 11 shows the measured off-axis PSFs and the result of a simulations using a remapping of the beam phase (Guyon 2003). With such a strong field aberration, measuring the focal plane plate scale is challenging and its value is a function of the metric used. In this article, we choose to adopt the noncoronagraphic PSF photocenter to measure plate scale. In Figure 10, the $4\lambda/D$ ring therefore shows where the photocenter of the PSF would be if the source was $4\lambda/D$ from the optical axis of the entrance telescope. The bright PSF core at this separation would be slightly outside the ring, but the fainter asymmetric diffraction arcs of the off-axis PSF would be inside the ring. The same photocenter metric is used to measure the angular sizes on the focal plane mask given in § 2.6. The plate scale on the coronagraph camera was measured by 2-D fitting of measured and simulated off-axis PSFs at several angular separations.

4.3. Coronagraphic Performance: Measurements

Coronagraphic performance is measured with the focal plane and Lyot masks in the beam and both the LOWFS loop and high-order wavefront control loop closed. Figure 12 (upper left) shows a raw image from the science camera. Most of the light reaching the camera is due to partial transmittance (at the 10^{-5} level) of the core of the focal plane mask, which produces a central peak in the image. The clear $1.65\lambda/D$ to $4.4\lambda/D$ opening in the focal plane mask is visible below this central peak.

Residual light is decomposed in two components by the wavefront sensing system, using the approach described in § 3.3:

1. An incoherent component composed of light that does not interfere with light extracted directly from the central PSF core. This component appears to be mostly stable in structure and

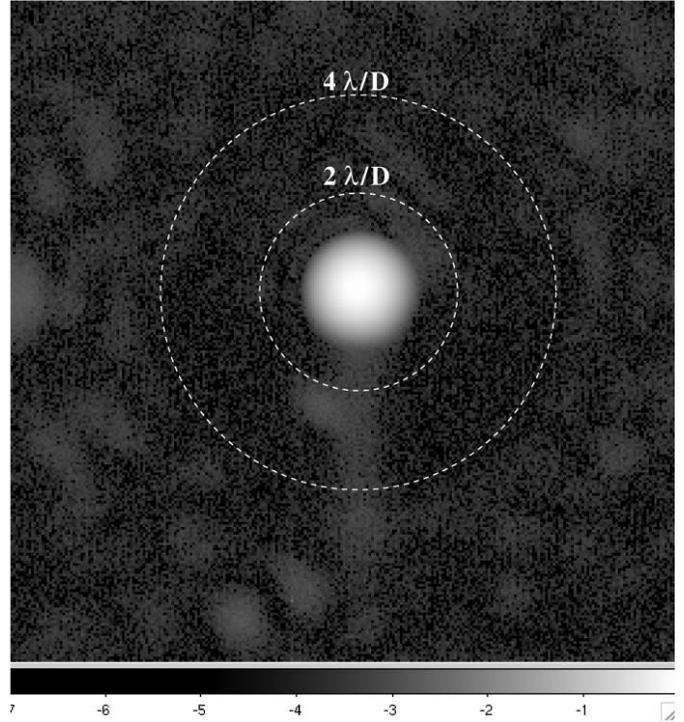


FIG. 10.—On-axis PSF without focal plane mask (log scale indicated at the bottom of the figure). The vertical bleeding feature extending downward of the PSF core is a camera artifact, and is removed when the focal plane mask is used. A faint ghost due to the entrance window of the camera is visible just beyond $2\lambda/D$.

varies from iteration to iteration between 1.5×10^{-7} and 2×10^{-7} in contrast.

2. A coherent component that is used to compute the DM shape for the next iteration. This component is at approximately 5×10^{-8} contrast and is decorrelated on timescales above the response time of the wavefront control loop. The estimation of this component varies greatly from iteration to iteration, and is sometimes below 10^{-8} contrast.

We also measured the χ^2 of the solution, and found it to correspond to a change in coherent light at the 4.5×10^{-8} contrast level during the measurement sequence, which is 10 s long. The M_{resp} optimization loop described in § 3.3.5 gave no noticeable improvement in χ^2 , suggesting that the initial calibration led to a good estimate of M_{resp} and that the observed χ^2 is indeed dominated by fluctuations in coherent light at the 4.5×10^{-8} contrast level.

Visual inspection of the coherent and incoherent portions of the light strongly suggests that this decomposition was successful. First, the images obtained show very little high spatial frequency noise. Although the reconstruction is performed independently for each pixel, both the coherent and incoherent components show speckles and features covering several pixels in these oversampled images. Second, the ghost to the lower left of the optical axis was properly analyzed by our algorithm.

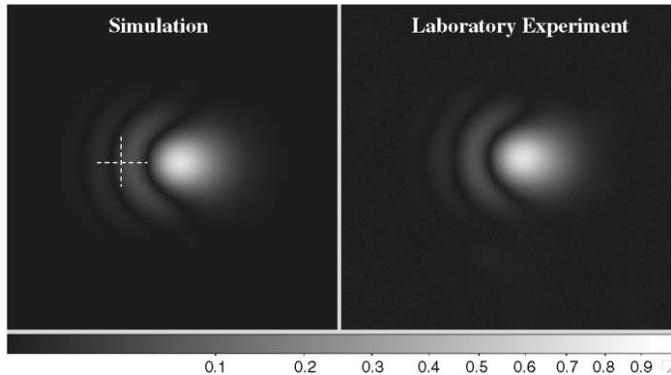


FIG. 11.—Off-axis PSF without focal plane mask, for a $1.7\lambda/D$ off-axis distance. A *cross* indicates the optical axis in the simulated image (*left*). Both images are shown to the same brightness scale (*bottom*).

This ghost, also visible in Figure 10, is due to internal reflection in the window of the CCD camera. Because the two surfaces of the window are not perfectly parallel, interference fringes can be seen across this ghost. Although the fringes are coherent, they are too narrow for the wavefront control system to remove. They are the only known feature in the image that (1) should be seen in the coherent image, (2) should not be seen in the incoherent image, and (3) cannot be removed by the wavefront control system. As shown in Figure 12, our coherent/incoherent analysis correctly identifies these fringe as coherent, and only a minute fraction of the fringes is present in the incoherent portion of the image (likely due to small motion of the fringes during the measurement sequence).

4.4. Analysis

The analysis results are summarized in Table 1 and discussed in this section.

4.4.1. Incoherent Light Component

The raw contrast is dominated by a very stable incoherent component at 1.63×10^{-7} , which is likely due to ghosts and/or polarization mismatch. Given the high number of air-glass surfaces (20), including some which are not antireflection coated, a ghost at this level would not be too surprising. A possible source of polarization mismatch is the ring apodizer, which has many $0.8 \mu\text{m}$ wide annular rings. The variations observed in the estimate of the incoherent light are due to the coherent light variation during the measurement, which affects the incoherent estimate. The data obtained is compatible with a fully static incoherent background, as would be expected from ghosts and polarization effects. We note that a fast varying coherent component could also be responsible for this light if it is varying sufficiently rapidly to appear incoherent within ≈ 1 s.

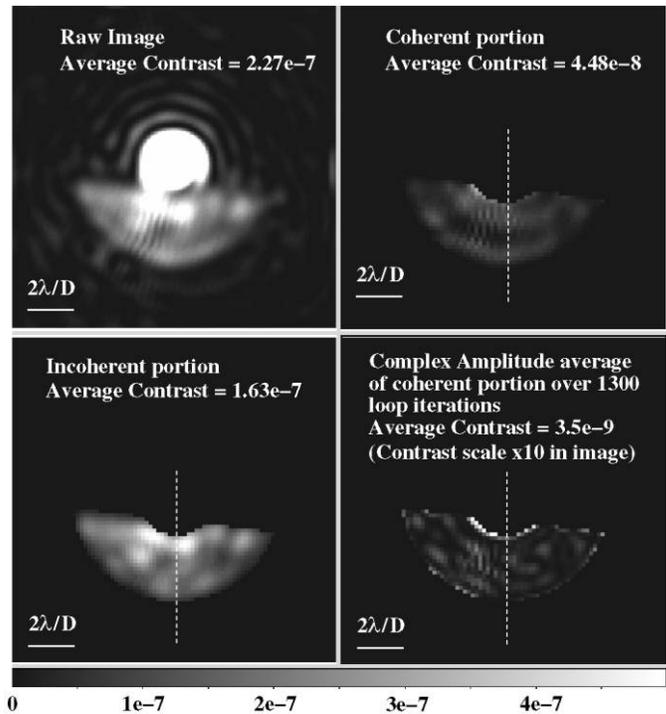


FIG. 12.—*Top left*: Raw coronagraphic image. A decomposition of the scattered light into a coherent component (*top right*) and incoherent component (*bottom left*) shows that the raw contrast is dominated by incoherent light. The coherent bias over a long period of time is shown in the *bottom right*.

4.4.2. Coherent light component

During the measurement sequence, coherent light varies by $\approx 1.4 \times 10^{-8}$ in contrast due to turbulence or vibrations in the system, as shown by the χ^2 analysis. The coherent light leak estimate over a 10 s period is $\approx 5 \times 10^{-8}$, which is at the level expected from the 4.5×10^{-8} variations shown by the χ^2 analysis. The large variation, from iteration to iteration, observed in the coherent light residual is due to the turbulence/vibrations in the system. We note that the lucky iterations where the 10 s averaged coherent light is estimated at below 10^{-8} are artefacts of the time averaging during the measurement period: even during these lucky periods, the coherent light did vary by $\approx 4.5 \times 10^{-8}$ between the 10 frames of the measurement.

4.4.3. Static Wavefront Aberrations

The wavefront control loop successfully removes static coherent speckles. Over a 4 hr period of time, we have measured the static coherent speckles to be at the 3.5×10^{-9} contrast level. Except for a known ghost on the camera window, we could not detect any residual bias above this level in the residual time-averaged coherent light.

The quality of the DM calibration (better than 10% for small displacements) and the use of 10 probes chosen to avoid

TABLE 1
CONTRAST BUDGET (AVERAGED OVER 1.6 TO $4.4\lambda/D$)

Term	Likely Origin	Value	Calibration
Incoherent light, static	Optical ghost, polarization	1.63×10^{-7}	Incoherent portion of WFS data
Coherent, variable (in $t < 10$ s)	Turbulence, vibrations	4×10^{-8}	Residual from WFS reconstruction
Coherent residual (in 10 s)		5×10^{-8} (typical) 7×10^{-9} (best)	Coherent portion of WFS data
Coherent, static (in 4 hr)	Uncorrectable coherent light	$< 3.5 \times 10^{-9}$	Time-averaged coherent light

systematic bias ensures that within a single wavefront measurement, the coherent wavefront error is below $0.1 \times \sqrt{10} \approx 0.03$ times the DM probe amplitude. Since DM probes were chosen to be at the 10^{-7} contrast level, the statistical measurement error due to DM calibration errors is expected to be below the $10^{-7} \times 0.03^2 \approx 10^{-10}$ contrast level for each wavefront measurement. Decorrelation between wavefront measurements further reduces the effect of DM calibration errors on the static wavefront aberration measurement accuracy. DM calibration and other potential sources of errors are further reduced by the system response matrix optimization described in § 3.3.5. At the 3.5×10^{-9} contrast level, the measurement of static wavefront aberration is therefore robust.

5. CONCLUSION

The results obtained in this experiment are especially encouraging for ground-based coronagraphy. The 2×10^{-7} raw contrast we have achieved in the 1.65 to $4.4\lambda/D$ angular separation range already exceeds, by 2 orders of magnitude, the raw contrast that can be hoped for in even a theoretically ideal extreme AO system (Guyon 2005). We note that with a more careful optical design and antireflection-coated optics, our experiment could probably have reached 5×10^{-8} raw contrast (level of coherent light leak in the current experiment). More importantly, we have demonstrated that with the coronagraph + wavefront control architecture adopted in our experiment, static coherent speckles can be pushed down very low (3.5×10^{-9}) in long exposures. Our system successfully removed long-term correlations in the coherent speckles, and their averaged level in long exposure was reduced with a $1/\sqrt{T}$ law. The combination of a high-performance PIAA coronagraph and a focal-plane based wavefront control therefore appears extremely attractive for ground-based extreme AO. In that regard, our experiment has been a successful validation of the key technologies and control algorithms of the Subaru Coronagraphic Extreme AO (SCEXAO) system currently in assembly. The major differences between the SCEXAO PIAA coronagraph and our laboratory prototype are (1) the need to design and operate a PIAA coronagraph on a centrally obscured pupil with thick spider vanes and (2) the need for corrective optics to recover a wide field of view. These two requirements have been validated in a separate

laboratory experiment using the final SCEXAO coronagraph optics (Lozi et al. 2009).

Our experiment was limited at the 2×10^{-7} contrast by an optical ghost and at the 5×10^{-8} contrast by turbulence or vibrations. The PIAA coronagraph could therefore not be tested to the contrast level required for direct imaging of Earthlike planets from space (approximately 10^{-10}), although several key concepts were demonstrated, including simultaneous operation of a low-order wavefront sensor using starlight in the PSF core and high-order wavefront sensor using scattered light in the science focal plane. New calibration schemes which will be very useful for high-contrast coronagraphy were also developed and validated, such as the system response matrix optimization loop, which can be executed as a background task to fine-tune the system.

PIAA coronagraph technologies for high-contrast space applications are now being actively developed and tested at NASA Ames Research Center and NASA Jet Propulsion Laboratory. A new set of PIAA mirrors was recently manufactured to higher surface accuracy than the ones used for this demonstration, and is being integrated within the High-Contrast Imaging Testbed (HCIT) vacuum chamber at NASA Jet Propulsion Laboratory. We note that the HCIT chamber has already demonstrated stability to the 10^{-10} contrast with a Lyot-type coronagraph (Trauger & Traub 2007). The experiment described in this article served as a precursor to this new step, which is aimed at reaching higher contrast (minimum goal of 10^{-9}) in broadband light using a two-DM wavefront correction. In parallel to this effort, a highly flexible high-stability air testbed at NASA Ames Research Center is coming online to further explore technology and architecture trades for PIAA systems.

In addition to pushing the contrast further, future laboratory demonstration of the PIAA coronagraph will explore chromaticity and wavefront control issues unique to a PIAA coronagraph architecture. The monochromatic experiment described in this article did not address these important points, and should therefore be considered as only a first step toward validation of the PIAA coronagraph technique for high contrast space-based exoplanet imaging mission.

This research was conducted with funding from the NASA Jet Propulsion Laboratory (JPL) and the National Astronomical Observatory of Japan. Technical input and advice from the members

of NASA JPL's High-Contrast Imaging Testbed (HCIT) team, NASA Ames Research Center's coronagraph team, and Princeton University's coronagraph team has been of considerable help in conducting this work, both for design/simulations and laboratory

implementation. In addition to providing laboratory space and infrastructure, Subaru Telescope made this research possible through major contributions by its technical staff (electronics, hardware, software).

REFERENCES

- Belikov, R., Kasdin, N. J., & Vanderbei, R. J. 2006, *ApJ*, 652, 833
- Galicher, R., Guyon, O., Otsubo, M., Suto, H., & Ridgway, S. T. 2005, *PASP*, 117, 411
- Givéon, A., et al. 2007, *Proc. SPIE*, 6691
- Guyon, O. 2003, *A&A*, 404, 379
- . 2005, *A&A*, 629, 592
- Guyon, O., Pluzhnik, E. A., Galicher, R., Martinache, F., Ridgway, S. T., & Woodruff, R. A. 2005, *ApJ*, 622, 744
- Guyon, O., Pluzhnik, E. A., Kuchner, M. J., Collins, B., & Ridgway, S. T. 2006, *ApJ*, 167, 81
- Guyon, O., Matsuo, T., & Angel, R. J. 2009, *PASP*, in press
- Lozi, J., Martinache, F., & Guyon, O. 2009, *PASP*, submitted
- Martinache, F., & Guyon, O. 2006, *EAS Publ. Ser.*, 22, 281
- Martinache, F., Guyon, O., Pluzhnik, E. A., Galicher, R., & Ridgway, S. T. 2006, *ApJ*, 639, 1129
- Pluzhnik, E. A., Guyon, O., Ridgway, S. T., Martinache, F., Woodruff, R. A., Blain, C., & Galicher, R. 2006, *ApJ*, 644, 1246
- Traub, W. A., & Vanderbei, R. J. 2003, *ApJ*, 559, 695
- Trauger, J. T., & Traub, W. A. 2007, *Nature*, 446, 771
- Vanderbei, R. J., & Traub, W. A. 2005, *ApJ*, 626, 1079
- Vanderbei, R. J. 2006, *ApJ*, 636, 528