

Exoplanet Imaging with the Giant Magellan Telescope

Johanan L. Codona

Steward Observatory, University of Arizona, Tucson, AZ, USA 85721

ABSTRACT

The proposed Giant Magellan Telescope (GMT) has a number of features that are well-suited to the task of imaging extra-solar planets in nearby star systems. The principal aid to this task is the large clear aperture segments which are relatively easy to apodize. This paper considers the methods currently envisaged to be of practical use for the task. In addition to star and planet fluxes, exoplanet imaging is dependent on aperture, throughput, bandwidth, beamwidth (FWHM), Strehl ratio (SR), and halo structure. Adaptive optics systems increase the SR, simultaneously dropping the residual scattered halo. This reveals the diffracted halo, which now becomes the limiting factor. Apodization reduces the diffracted halo, but at a cost in terms of throughput and a corresponding increase in photon noise. Since the best known ideal apodizations also have very low throughputs, they are not the best choices for ground-based exoplanet imaging. In addition, the ultra-low diffracted halos from these apodizations provide no benefit below the residual scattered halo, which is not helped by apodization. We consider instead a family of apodizations that have sufficiently dark diffracted halos, while retaining relatively high throughputs. One form of apodization can be applied to the GMT pupil using replicated apodization of individual segments, providing a low-halo survey mode that is high throughput and matched to the AO system. Since the reduced halo from the apodized segments only allows high-contrast detection to within a few $\lambda/D_{segment}$ of the star, the single segment methods are limited by the segment size. We also consider the potential for apodizing the full aperture for high contrast at a few times λ/D_{full} through the use of an applied phase pattern, using either the adaptive secondary or a separate phase mask. We conclude that the phase mask method offers the best advantage for S/N since it does not lose light like the apodization schemes. However, it does have a restricted azimuthal search area, requiring multiple exposures to complete a survey. It appears to be the clearly best method for examining the exoplanet once discovered. It should be possible to apply offsets to the GMT's adaptive secondary to achieve contrasts of 10^{-5} at $2 \times \text{FWHM}$ (27 mas at $1.65 \mu\text{m}$, 80 mas at $5 \mu\text{m}$) and 10^{-6} at $3 \times \text{FWHM}$ (42 mas at $1.65 \mu\text{m}$, 200 mas at $5 \mu\text{m}$).

Keywords: large telescopes, extreme adaptive optics, high-contrast imaging, extrasolar planets.

1. INTRODUCTION

The Giant Magellan Telescope (GMT) is a proposed multi-segment telescope consisting of seven 8.4 m circular mirror segments¹ (figure 1). The outermost diameter of the pupil is over 26 m, equivalent in light gathering area to a circular pupil of diameter 21.9 m. The outer six segments are unobstructed and have relatively few degrees of freedom to control for joint phasing. This paper considers the use of the GMT for imaging extrasolar planets in nearby star systems. The principal aid to this task is the large clear aperture segments which are relatively easy to apodize by shading and applying phase profiles using the adaptive secondary.

In addition to the relative star and planet fluxes and separation, exoplanet imaging is dependent upon the detailed PSF of the telescope, and the ability of the AO system to reduce both the speckles and their temporal correlations so as to aid in averaging. For the purposes of this paper, we will assume that the residual AO halo is dominated by fitting error, and that speckle noise has been controlled through the use of advanced future methods and algorithms.




Figure 1. The Giant Magellan Telescope (GMT).




Figure 2. GMT pupil configurations.




Figure 3. GMT7 and GMT6 PSFs computed for a wavelength of $1.65\mu\text{m}$ and a 10% bandwidth.

The residual fitting error halo contributes a flat background in the immediate vicinity of the star, and its photon noise cannot otherwise be reduced. If the pupil diffraction pattern exceeds the AO halo, the photon noise will be dominated by it rather than the AO residual. Since there are a variety of techniques such as apodization that can be applied to reduce the

diffraction pattern, it makes sense to apply them to reduce the diffraction to a level below the AO halo such that detection is again limited by scattering rather than diffraction. This places the system performance constraint on the system element that is the most expensive and difficult to control, giving the best achievable performance. Any residual unaveraged speckles or deterministic structure in the PSF can be removed by differencing methods², leaving a detection limit not far in excess of the photon noise.

In the following sections, we consider the expected contrast ratio for the residual AO halo, setting the level below which we must drop the diffraction pattern of the telescope. We then consider apodization methods, preserving the effect of the large aperture by apodizing the individual circular outer segments. We use an aggressive apodization that is not far from the optimal performance, and try a high-throughput band-stop apodization as an alternative. We then try another method that uses a deformable mirror (DM) or phase plate to explicitly cancel the halo in a search region around the star. This method has the advantage that it does not affect the throughput of the telescope, but it does scatter light to create the anti-halo, which removes light from the exoplanet image. Finally, the different methods are compared.

Contrast of AO and diffraction halos

For wavelengths shorter than those where the sky background dominates, the stellar halo is the primary background against which we will attempt exoplanet detection. This halo consists of two principal components: a diffraction pattern from the telescope, and a random residual halo caused by various inaccuracies in the AO system. For feasible Strehl ratios, the contrast ratio between the diffraction-limited core of the PSF and the residual AO halo is about five decades at best. We can see this by estimating the core/AO-halo contrast ratio by energy accounting, rapidly leading us to the contrast between the diffraction-limited core and the flat fitting-error AO halo as

$$\frac{(1-S)}{S} \frac{1}{fN_{actuators}}$$


where S is the Strehl ratio, f is the fraction of the diffraction pattern energy contained in the central peak, and $N_{actuators}$ is the number of actuators, currently expected to be approximately 4700. The GMT seven-mirror configuration has an f of 67%. For a 50% Strehl ratio, the contrast ratio will be approximately 3×10^{-4} , not reaching a contrast of 10^{-5} until the Strehl ratio hits 97%. Thus, the practical expectation is that the contrast will be in the vicinity of 10^{-4} – 10^{-5} , only dropping below that level at longer wavelengths where the sky background becomes the constraint. For this paper, we will use 10^{-4} as the nominal background against which the detections are made. The detection of exoplanets, whether made against the residual AO halo or the diffraction pattern of the PSF, are both ultimately limited by photon noise. The two cases are very different in the cost and complexity of reducing their effect. Since reducing the AO halo is the far greater challenge, we will concentrate on reducing the diffraction pattern such that scattering again becomes the limit to detection. We will find out through calculation how much this improves the sensitivity.

2. THE GMT PUPIL AND PSFS


The GMT pupil consists of seven 8.4 m mirror segments, arranged to form segments of a single 18 m focal length paraboloid. The central mirror is currently planned to have a 3.1 m central obstruction for the adaptive secondary, while the ring of six outer segments are essentially unobstructed parabolic segments. Since the circular mirror segments in the outer ring are aligned with their upper edges tangent to the paraboloid, they project elliptical outlines in the pupil projected along the telescope axis. The outside mirrors are placed at a radius from the telescope axis to give a separation of 0.25 m between mirrors. The adaptive secondary is a Gregorian design, with a segmented adaptive secondary. For the sake of analysis, we will consider the actuator spacing to be the projected equivalent of the LBT, with 672 actuators for each mirror. The resulting pupil consists of an obstructed circular segment with a ring of six elliptical segments (figure 2a). Another useful configuration is the six-mirror configuration that results when the central mirror is blocked (figure 2b). This configuration consists of a more symmetric arrangement of unobstructed elliptical segments, which is more convenient for some of the apodization configurations described below.

3. INDIVIDUAL SEGMENT APODIZATION


The contrast ratio for the AO halo is typically going to be of order 10^{-4} – 10^{-5} . For the photon noise to be set by adaptive optics rather than diffraction, we should modify the diffraction pattern such that it is well below the AO halo. A requirement of 10^{-5} – 10^{-6} suffices for most cases. The traditional approach to this is through pupil apodization. To achieve high



(a) Chebyshev 5, 6, and 7 decade shading profiles.



(b) Chebyshev 7-decade mask applied to outer segments.



(c) Chebyshev 6-decade PSF.

Figure 4. Chebyshev window functions applied to the outer six segments.

contrast diffraction patterns at as close of spacings as possible, we should really jointly apodize across all of the segments in a coordinated way. The GMT, with its relatively small number of large segments makes this an attractive possibility. However, for our initial apodization studies, we have adopted a straightforward variant of the classic apodizations by treating each segment in isolation. This makes use of symmetric radial shading profiles in the individual segments while preserving much of the resolution of the overall pupil by the placement of the segments. This technique works well, but is constrained in that the inner working distance is limited to some multiple of $\lambda/D_{segment}$ instead of a multiple of λ/D_{full} . For many situations however, this is still of use.




Figure 5. High-throughput (high- τ) apodized GMT6 and resulting PSF.

For this discussion, we will only consider apodizations of the unobstructed circular apertures. The center mirror with its secondary obstruction and spider supports cause a number of problems that we are trying to avoid, so we chose to block it out entirely, using the GMT6 as the base configuration. Also, even though it would be possible to adapt the symmetric radial apodizations to the elliptical projected outlines of the outer segments, it is only a 3% drop in throughput to use a masked inscribed circular outline for the segment apodizations. Therefore, we examined the case of a ring of six circular apodized segments.

Chebyshev Apodization

The optimum apodization of an unobstructed circular aperture was derived by Slepian³ and is given by the so-called generalized prolate spheroidal function. This function is not readily available however, since it is quite complicated to compute. Instead, we found that the ubiquitously available Chebyshev profile⁴ was an excellent replacement. It is easy to use, has nearly the same characteristics as the ideal function, and allows easy setting of the contrast ratio of the halo in the search area. We only made one simplifying modification in that the Chebyshev profile has a bright edge in the outermost pixel of the digital mask. We removed this outermost pixel and renormalized the transmission so the mask was completely transparent in the center of the circular segment. The Chebyshev mask then gives a radial profile which was then interpolated onto the 2-D array for the PSF calculations. The Chebyshev low-pass filter stop-band factor was found to adequately give the diffraction halo contrast ratio. We report here the result for the cases of 5, 6, and 7 decades of rejection (labeled “cheby5-7”). As in the ideal Slepian case, the throughput drops as the working area halo is pushed down, and the FWHM of the segment PSF increases. Fortunately in our case, the FWHM is not determined by the segment PSF, but by the PSF of the ring of six apodized segments. This keeps the FWHM to a value only slightly greater than that of the unapodized pupil. However, FWHM is a metric that is of little use when you are interested in extremely faint objects in the halo. A more practical metric is the radius at which the average PSF reaches a level of say 10^{-5} or 10^{-6} of the peak.

High-throughput Apodization

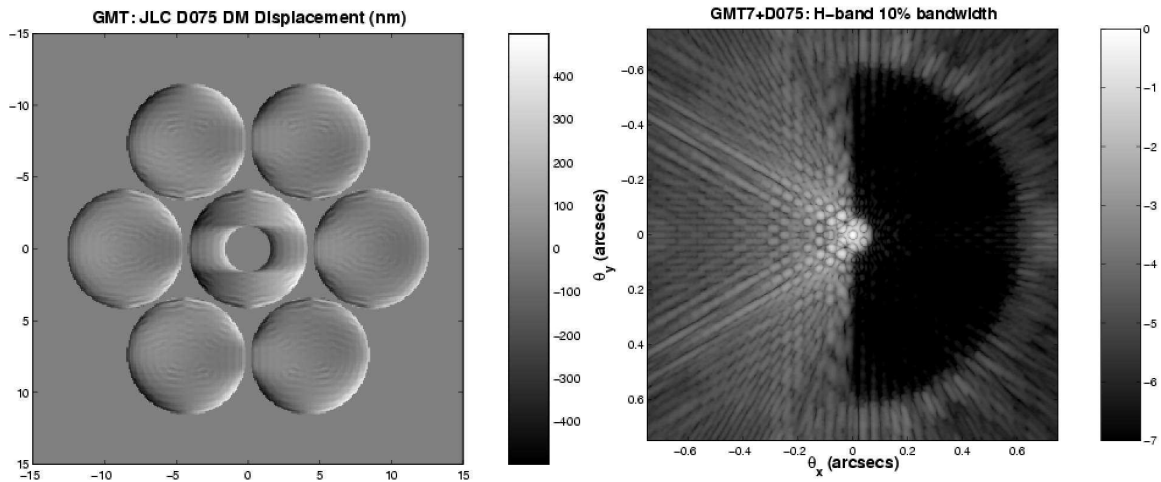
The most serious problem with the above apodization is the loss of light, which in the circular case can easily give throughputs of less than 10%. In the simplest terms, this problem is due to the fact that the ideal low-pass filter type apodizations have their throughput maximum at the center of the aperture and thus they have a throughput maximum over only a small area of the pupil. This suggests that any extra constraint that can introduce oscillations in the radial apodization profile could be adjusted such that the profile did not reach its maximum in the center, but at some larger radius. The resulting

throughput would be correspondingly greater. This problem can be treated by introducing both an inner and an outer working radius. There were a number of details in performing this calculation that are beyond the scope of this paper, but do not affect the conclusions for this application. The shape of the computed apodization as well as the throughput and the average halo stop band level were sensitive functions of the inner and outer working radii compared with λ/D . Therefore, a brute force calculation of all inner and outer radii within a practical set of ranges allowed the calculation of a set of metric maps that allowed selection of high throughput or high sensitivity apodizations. We chose a profile that had a relatively high throughput, a fairly aggressive inner working distance, and a useably wide stop band for the example here. The profile and its corresponding PSF is shown in figure 5. The resulting PSF differs from the Chebyshev case in two significant ways: the throughput is greater, and there is an obvious increase in the halo level beyond the outer working distance. For searching for objects within the halo stop band, the higher throughput and greater encircled energy within the main beam will allow a given S/N to be reached in a much shorter time (about a third of the time compared to Cheby6). The relative performance is summarized in table 1. Since $S/N \propto \sqrt{T}$, the required exposure time is proportional to the inverse square root of the S/N metric. Therefore, the high- τ apodization will detect a given exoplanet in roughly 1/2, 1/3, or 1/5 of the time compared to Cheby5, Cheby6, and Cheby7 respectively. This is with a FWHM at $1.65 \mu\text{m}$ of 13.7 mas, and drops to a 6 decade contrast by 82 mas from the star. The Chebyshev apodizations maintain essentially the same FWHM of 13.9 mas, while slowly increasing the high-contrast working radius and decreasing the throughput with increasing halo rejection. As long as the search area is beyond about $5\text{--}7 \lambda/D_{full}$, any of these apodizations are reasonable choices.

4. PHASE MASKS AND DIFFRACTED ANTI-HALOS


The problem of diffracted halo control by segment apodization is loss of light. Because of the segment replication, the FWHM is not significantly increased, even though the the full width at lower levels is. Another possibility is to use the method described in Codona and Angel⁵, that uses a DM or a specially constructed phase plate to create a matched anti-halo for the diffraction pattern for a particular detection band. The method is chromatic and degrades with bandwidth, but since our goal is merely to push the diffracted halo below the level of the residual AO halo, the method is good enough to work over common detection bands. The idea is to superpose small-amplitude sinusoidal ripples in the DM to diffract some of the core starlight to counteract speckles in the focal plane. The orientation and spatial frequency of the sinusoid places the anti-speckle, while the sinusoid's phase on the optical axis determines the phase of the anti-speckles at the design wavelength. The resulting anti-speckle adds coherently with the original speckle reducing its amplitude. Of course, this procedure has a number of side effects, like introducing undesired structure in the direction diametrically opposite to the star, since small amplitude variations create a halo that is anti-Hermitian, while the diffraction pattern due to the pupil is Hermitian. The anti-halo is also a replica of the PSF, which may have structure that is not present in the halo to be controlled. Finally, the sinusoidal variations in the pupil plane create diffraction orders that appear at multiples of the speckle distance from the star. These problems are all reasonable though, and a diffraction mask can be computed by iterating the method.

The result is a spatial displacement pattern that can be applied to a DM, or introduced via a separate optic. The cleared area in the focal plane is subject to a number of constraints. The controlled focal plane area divided by the typical speckle area, $(\lambda/D_{full})^2$, must be less than the number of actuators, and the total halo power in the controlled region should be significantly less than power in the core. If this isn't true, the exoplanet will be distorted as well as having significantly fewer core photons, increasing the time required to make a detection. Finally, the obvious point that the outer edge of the search region must be inside the control radius of the DM. Since in this application the desire is to control a static diffraction pattern, the control radius restriction can be lifted by using a specially manufactured phase plate to apply the phase. The control radius with the proposed GMT actuator density is approximately $45\lambda/D_{full}$ but is actually dependent on the details of the direction from the star. Figure 6 shows a case which falls far inside the control radius ($6\lambda/D_{full}$) and could feasibly be implemented using the AO DM, while figure 7 shows a case where the outer radius exceeds the capabilities of the DM ($60\lambda/D_{full}$) and would have to be implemented using a separate optical element. The mirror displacement for the smaller sector is 0.53λ (half the total phase shift) while the larger sector requires a maximum displacement of 0.97λ . At $1.65\mu\text{m}$ this amounts to 0.88 and $1.6\mu\text{m}$ respectively. With the anticipated extended performance of the adaptive secondary, the increased throw is not expected to be an issue. The symmetry consideration means that it is not possible to control the diffracted halo using the phase mask method over a region subtending more than 180° azimuthally about the star. The total power restriction also keeps us from simultaneously clearing a large area and moving extremely close to the star. However, this restriction appears to be very forgiving, and we have found that search areas can be brought in extremely close to the star, even comparable to $2\lambda/D_{full}$. The parameters for the two phase mask cases are summarized in table 1.



(a) DM displacement required to build a 180 degree anti-halo.

(b) Resulting PSF with a 10 percent bandwidth.



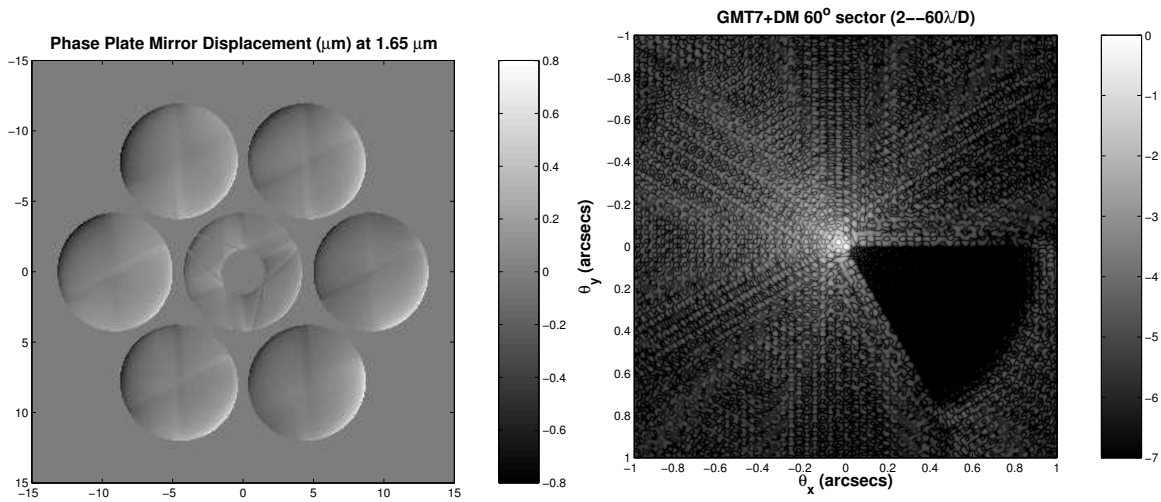
(c) Less extreme phase mask to create a small 60 degree anti-halo sector from 2 to 6 lambda/D.

(d) Resulting PSF.

Figure 6. Examples of GMT7 phase plate solutions creating dark regions in the halo.

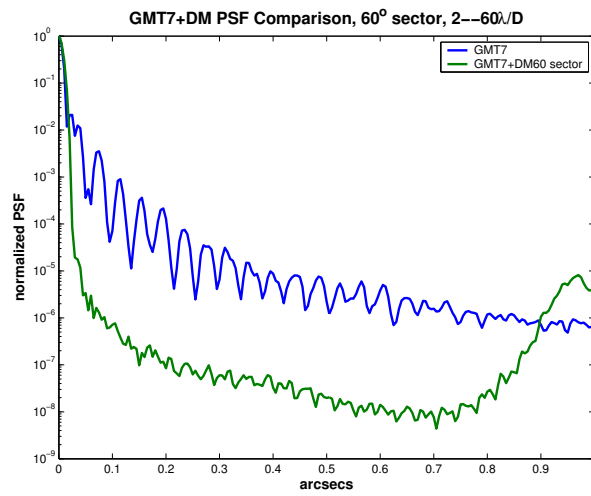
5. PERFORMANCE ESTIMATES

The performance numbers for the tested cases are summarized in table 1. The S/N for an AO halo-limited detection is proportional to $f\sqrt{\tau}/FWHM$, where f is the fraction of the encircled energy in the PSF core, and τ is the intensity throughput. Detection in the unapodized GMT7 and GMT6 cases would be limited by their diffraction patterns, and so the S/N metric is not accurate. In those cases, the typical close-in contrast is more like 10^{-2} than 10^{-4} , with the result that their metric values should be reduced by a factor of roughly 10 (i.e. $\sqrt{\text{diffraction}/AO}$). Note that, as expected, as the Chebyshev apodized halo drops farther below the AO halo, the sensitivity drops. The high-throughput ring apodization



(a) Pupil mirror displacement required to build a 60 degree anti-halo from 2 to 60 lambda/D.

(b) Resulting PSF, displayed on a seven-decade logarithmic scale.



(c) Narrowband comparison of the radial average PSF for the 60 degree sector phase mask, and the unmodified GMT7.

Figure 7. GMT 7 Anti-halo phase plate correcting a 60 degree sector of the focal plane.

Parameter	GMT7	GMT6	Cheby5	Cheby6	Cheby7	High- τ	DM2-6	DM2-60
Intensity-weighted area (m ²)	368.6	323.1	58.1	47.9	40.8	90.2	368.6	368.6
Throughput re clear (%)	100	100	18.0	14.9	12.6	24.5	100	100
Throughput re GMT7 (%)	100	87.7	15.8	13.0	11.1	21.5	100	100
FWHM @ 1.65 μ m (mas)	14.4	13.5	13.8	13.9	13.9	13.7	14.8	14.4
Core encircled energy, f (%)	67.5	50.9	30.1	26.6	23.4	36.4	39.5	23.5
Radius @ 5 decades (mas)	250	490	80	97	104	78	27	27
Radius @ 6 decades (mas)	400	830	225	102	113	82	42	65
S/N metric ($f\sqrt{\tau_{GMT7}}/FWHM$)	46.9*	35.2*	8.7	6.9	5.6	12.3	26.8	16.3

* These cases are not AO halo limited. They should be reduced by the square root of the actual contrast. Estimate them as approximately 10 \times smaller. PSF radii are computed for 1.65 μ m and the throughput used for comparison in the metric is computed relative to the clear GMT7 case.

Table 1. Performance numbers for the various configurations.

performs comparably to the five-decade Chebyshev apodization, but with a more limited search area. The clear standout is the phase plate or DM method, which more than doubles or triples the sensitivity. An additional advantage the phase method has over the other methods described here is that the inner working distance can be brought in extremely close to the star. It is conceivable that similar close-in contrasts could be achieved using joint intensity apodization across all of the segments, but with the inevitable loss of light, the S/N metric will likely remain considerably lower.

Table 1 leads us to a number of important conclusions. First, even though intensity apodization techniques lose a great deal of light, the individual segment apodization method allows us to preserve most of the angular resolution of the telescope. This results in a better improvement in the S/N metric than would be expected if apodizing a single-segment telescope. Second, even though the half-max beamwidth does not change much with the various methods, the much more significant widths at 10⁻⁵ and 10⁻⁶ do. At 1.65 μ m, the unmodified GMT7 PSF does not fall below a contrast of below a contrast of 5 decades until we look beyond a quarter of an arcsecond; over 17 FWHM. All of the other methods do much better. The per-segment intensity apodization methods yield 5 decade contrasts at better than 6–7 FWHM, with expected gains even beyond that if we were to jointly apodize all the segments. Finally, since it was a joint method across all of the segments, the DM “phase apodization” method achieved the most rapid drop to high contrast, with 5 decades in less than 2 FWHM. What is the most remarkable is that this high contrast (at least in the more restricted search area case) should be possible by just biasing GMT’s adaptive secondary. This means that with the GMT’s highly-optimized infrared design and the additional use of the adaptive secondary to reduce close-in diffraction, it should be possible to obtain a contrast of 10⁻⁵ at 82 mas and 10⁻⁶ at 127 mas at 5 μ m. Integration and the addition of differential techniques can reasonably be expected to increase the detectable contrasts by another 2 – 3 decades. This will yield high sensitivity to giant exoplanets which are expected to be bright at 5 μ m.

6. DISCUSSION

In considering exoplanet imaging with the GMT, it is clear that the raw diffraction pattern will limit the detection of faint exoplanets near the star. We have considered apodization as a method for controlling the diffraction, as well as a phase apodization method to explicitly cancel the diffraction pattern. For useful infrared detection bands, at best we can expect the AO residual halo to be no more than 4 or 5 decades below the peak of the star’s PSF. If we do nothing to the pupil, the contrast close to the star varies around 2–3 decades down from the peak. The most interesting region for high-resolution imaging, close to the diffraction limit for the overall aperture, has a natural contrast of around two decades below the peak. Taking all of this into account through the S/N metric in table 1, and reducing the values for the GMT7 and GMT6 by a factor of 10 for their higher limiting halos, we see that the phase mask methods are the winners. The loss of light from the apodization methods nearly cancel out the gains made by reducing the diffraction halo. The only apodization cases that were reasonable improvements were the least extreme Chebyshev case (Cheby5) and the high-throughput mask. Because these were only applied to the individual outer segments, their advantages are not available for extremely close-in surveys. To make the advantage work to the extreme limit would require that the apodization be performed jointly across all of the segments. The biggest winner would appear to be the phase apodization method since it reaches an S/N 3–5 times greater

than the other methods in the same amount of time. However, the phase method has an inherently restrictive azimuthal search area, requiring different sectors around the star to be searched sequentially. Although we analyzed 60° sectors for this paper, sectors as large as 180° are possible. Since $S/N \propto \sqrt{time}$, the search time advantage is proportional to the square of the S/N metric divided by the number of required search sectors. Even for an S/N advantage of $3\times$ and 60° sectors, the advantage would still be $3^2/6 = 1.5$. More careful tailoring of the mask and search area should keep the advantage of using the phase mask method to at least two or three times that of using the unmodified GMT7 for searching very close to the star. Once an object has been discovered, the phase mask method has a clear advantage in that it can be tailored for the specific region of interest, allowing more detailed study.

7. ACKNOWLEDGEMENTS

This research is supported by the NSF under grant AST-0138347.

References

1. Johns, M., R. Angel, R. Bernstein, D. Fabricant, P. McCarthy, M. Phillips, S. Shectman, "Status of the Giant Magellan Telescope (GMT) project", SPIE **5489**-28, 2004.
2. Biller, B. A., L. Close, R. Lenzen, W. Brandner, D. McCarthy, E. Nielsen, M. Hartung, "An algorithm for the suppression of speckle noise for simultaneous differential extrasolar planet imaging at the VLT and MMT", SPIE **5490**-52, 2004.
3. Slepian, D., "Analytic solution of two apodization problems", JOSA, **55**(9),1110, 1965.
4. IEEE 1979, *Programs for Digital Signal Processing*. (New York: John Wiley & Sons), Program 5.2.
5. Codona, J. L. and Angel, R., "Imaging extrasolar planets by stellar halo suppression in separately-corrected color bands", ApJ **604** L117, 2004.