The Astrophysical Journal, 646:1260–1274, 2006 August 1
© 2006. The American Astronomical Society. All rights reserved. Printed in U.S.A.

LINEAR AND BAYESIAN PLANET DETECTION ALGORITHMS FOR THE *TERRESTRIAL PLANET FINDER*

N. Jeremy Kasdin
Mechanical and Aerospace Engineering, Princeton University,
Princeton, NJ 08544; jkasdin@princeton.edu

and

Isabelle Braems
LEMHE / ICMMO, UMR 8182, Université Paris-Sud,
91405 Orsay Cedex, France; isabele.braems@lemhe.u-psud.fr
Received 2005 December 16; accepted 2006 April 5

ABSTRACT

Current plans call for the first *Terrestrial Planet Finder* mission, *TPF-C*, to be a monolithic space telescope with a coronagraph for achieving high contrast. The coronagraph removes the diffracted starlight allowing the nearby planet to be detected. In this paper, we present a model of the planet measurement and noise statistics. We use this model to develop two planet detection algorithms, one based on matched filtering of the point-spread function (PSF) and one using Bayesian techniques. These models are used to formulate integration time estimates for a planet detection with desired small probabilities of false alarms and missed detections.

Subject headings: methods: data analysis — planetary systems

1. INTRODUCTION

In 2004, NASA announced that it plans to launch the *Terrestrial Planet Finder–Coronagraph* (*TPF-C*), a space-based visible light telescope equipped with a coronagraph to detect earthlike planets around other stars, by 2016. Work is proceeding on developing a baseline design that would meet the scientific requirements. This currently consists of a 8 m × 3.5 m off-axis Cassegrain optical telescope combined with a starlight suppression system (SSS) to achieve the needed starlight rejection. The SSS is equipped to use either conventional Lyot coronagraphs or pupil masks (shaped or apodized). At Princeton, we have been studying what we call "shaped pupil coronagraphs" as an alternative to the traditional Lyot design (Kasdin et al. 2003, 2005; Vanderbei et al. 2003a, 2003b, 2004). These are coronagraphs with apodized entrance pupils that rely solely on one/zero binary openings, resulting in a more manufacturable and robust design.

Critical to *TPF-C*'s success is the ability to survey as many stars as possible and maximize the chances of detecting potential planets. To do this, a thorough understanding of the data analysis is required to obtain reasonable estimates of the integration time for each observation. Brown (2005) discusses at length the question of survey completeness and its relation to integration time. In this paper, we focus on a more careful and complete model of the measurement and noise statistics and develop a probabilistic detection criterion. We also analyze a Bayesian detection approach to see if detection time improvements can be made. These expressions can be used for further completeness studies and estimates of the efficacy of the *TPF-C* search program. We also note that these techniques apply to the photometry and detection of any dim object, not just planets in the context of *TPF-C*.

2. POINT-SPREAD FUNCTION

Before describing the detection approach, in this section we develop the needed definitions, units, and notation for photometry. We assume a perfect imaging system and scalar Fraunhofer theory such that the electric field at the image plane of the camera is the Fourier transform of the field entering the camera at the exit pupil of the starlight suppression system (SSS),

$$E_f(u,v) = \frac{e^{j(k/2f)\left(u^2+v^2\right)}}{j\lambda f}\int\int_S E_o(x,y)e^{-j(2\pi/\lambda f)(xu+yv)}\,dx\,dy, \tag{1}$$

where E_o is the electric field at the exit pupil aperture, f is the focal length of the camera, λ is the wavelength of the light, and the set S represents the exit pupil aperture, which, for an open square aperture, as an example, is given by

$$S = \left\{(x,y) : -\frac{D}{2} \le x \le \frac{D}{2}, -\frac{D}{2} \le y \le \frac{D}{2}\right\},$$

where D is the width of the aperture.

The optical system prior to the exit pupil of the SSS consists of various reflections and correction elements all of which have a certain efficiency. If we take $E_i(\lambda)$ to be the uniform electric field at the entrance pupil, then E_o can be written

$$E_o(X,Y,\lambda) = \eta\mathcal{A}(X,Y)E_i(\lambda),$$

where $\eta < 1$ is the overall attenuation of the optical elements up to the exit pupil (including any reduction due to an image plane mask), $\mathcal{A}(X,Y)$ is the apodization of the exit pupil, and we have explicitly shown the dependence on wavelength. For a classical or band-limited coronagraph, $\mathcal{A}(X,Y)$ is one/zero valued and just the Lyot stop shape. For a shaped pupil coronagraph, $\mathcal{A}(X,Y)$ represents the entire starlight suppression and is also one/zero valued.

For a point source, the image plane detector measures the irradiance of the arriving field, $I_f = (c\epsilon_o/2h)|E_f|^2$, in photons s^{-1} m^{-2} μm^{-1}, where c denotes the speed of light, h is Planck's constant, and ϵ_o is the permittivity of free space. Squaring the

magnitude of the electric field in equation (1) gives the irradiance at the detector,

$$I_f(\lambda,u,v) = \eta^2 \frac{c\epsilon_o}{2h}|E_i(\lambda)|^2 P(\lambda,u,v) = \eta^2 I_i(\lambda) P(\lambda,u,v), \quad (2)$$

where $I_i(\lambda) = (c\epsilon_o/2h)|E_i(\lambda)|^2$ is the irradiance of the point source at the reference wavelength and $P(u,v)$ is the dimensionless point-spread function (PSF) associated with the exit pupil of the SSS,

$$P(\lambda,u,v) = \frac{1}{\lambda^2 f^2}\left|\int\int_S \mathcal{A}(x,y) e^{-j(2\pi/\lambda f)(xu+yv)}\,dx\,dy\right|^2. \quad (3)$$

We note that the value of the integral in equation (3) at the origin is the square of the integral of the apodization, which for an open aperture (or shaped pupil or Lyot stop) is the square of the area of the exit pupil. We use A to represent the area of the entrance pupil of the telescope and A_c to represent the area of the exit pupil of the coronagraph (Lyot stop or shaped pupil). For an open telescope with no coronagraph, $A_c = A$. For apodized pupils, we call A_c the "pseudoarea" as it includes the reduction in throughput due to the continuous apodization.

There are three expected sources that produce electric field entering the telescope: the parent star [$I_s = (c\epsilon_o/2h)|E_s(\lambda)|^2$], the planet [$I_p = (c\epsilon_o/2h)|E_p(\lambda)|^2$], and the exozodiacal dust. While the first two sources are point sources, the exozodi is a uniform extended source with surface brightness $\mathcal{F}_{\rm ez}(\lambda)$ in units of photons s^{-1} m^{-2} μm^{-1} sr^{-1}.

For a planet offset on the sky by the angles (θ_x,θ_y) the electric field at the entrance pupil becomes $E_p(\lambda)\exp[2\pi i(\theta_x x/\lambda + \theta_y y/\lambda)]$, and by the Fourier shift theorem the irradiance at the detector becomes

$$I_f(\lambda,u,v) = \eta^2 I_p(\lambda) P\big(\lambda, u-\theta_x f, v-\theta_y f\big). \quad (4)$$

For a diffraction-limited system, we can assume that the uniform exozodical flux also creates a uniform flux density on the image plane. A careful analysis would include a convolution of the exozodi with the point-spread function, which could be significant if the zodi were nonuniform. However, for this analysis we simplify by assuming a uniform zodi, and thus we can ignore the details of the PSF. The resulting irradiance at the detector is given by

$$I_{\rm ez} = \mathcal{F}_{\rm ez}(\lambda) P(0,0)\Omega_{ij}, \quad (5)$$

where Ω_{ij} is the solid angle on the sky subtended by a single pixel, in steradians.

3. MEASUREMENTS

In this section we develop a model for the detector measurements and corresponding noise statistics. This is essential for developing the detection algorithm and corresponding detection probabilities. We assume here that the wavefront control system has successfully reduced the aberrated light to a uniform background in the discovery region or that we are dominated by the exozodiacal background. Thus, the intensity at the planet location consists of the planet PSF from equation (4) plus a uniform background, $I_b = I_{\rm ss} + I_{\rm sc} + I_{\rm ez} + I_r$, where $I_{\rm ss}$ is due to the residual suppressed starlight, $I_{\rm sc}$ is due to scattered light, $I_{\rm ez}$ is the exozodiacal dust background, and I_r is the effective readout noise. All are in units of photons s^{-1} m^{-2} μm^{-1}. The exozodi term was defined in § 2. The suppressed starlight is given by the residual value of the PSF, after any coronagraph, at the planet location,

$$I_{\rm ss} = I_s(\lambda) P\big(\lambda, \theta_x f, \theta_y f\big). \quad (6)$$

We also assume that because of cosmic-ray concerns the eventual signal will consist of a stack of many frames read out at a faster rate, and thus it is sensible to treat the readout noise as absorbed into the uniform background. We also assume that the readout noise is small compared to the other sources of background noise. In fact, it is expected that the dominant background source will be the exozodiacal light, perhaps as much as 3–10 times the brightness of the planet.

3.1. *Poisson Statistics*

The photon arrival at each pixel is determined by a Poisson distribution. That is, the probability of a given measurement, $z_{ij}(t)$, at pixel (i,j) and time t is given by

$$\Pr\big(\mathbb{Z} = z_{ij}(t)\big) = e^{-I_{ij}t}\frac{(I_{ij}t)^{z_{ij}}}{z_{ij}!},$$

where I_{ij} is the mean arrival rate of the photons, or intensity, at pixel (i,j) in photons per second. Thus, if $z(t)$ is the actual measured number of counts at any pixel,

$$\mathcal{E}[z(t)] = It = \mu(t),$$
$$\mathcal{E}\Big\{[z(t)-\mu(t)]^2\Big\} = It = \sigma^2(t),$$

where I is the intensity at that pixel.

The photon arrival rate, or intensity, in photons per second, at pixel (i,j) is the integrated irradiance over the pixel area and a reference passband about a reference wavelength λ_r, which for the planet is given by

$$I_{ij} = \eta^2 I_p(\lambda_r)\int_{\lambda_r-\Delta\lambda/2}^{\lambda_r+\Delta\lambda/2}\int\int_{\Delta\alpha_{ij}} P\big(\lambda, u-\theta_x f, v-\theta_y f\big)\,d\lambda\,du\,dv, \quad (7)$$

where $\Delta\alpha_{ij}$ represents the region covered by pixel (i,j) and $\Delta\lambda$ is the reference wavelength band given by the largest band of interest.[1] In general there is not much more that can be done without a specific point-spread function (and then these integrals can be difficult in closed form). However, since we are seeking a simple and conservative metric, we can obtain a useful result by assuming that the bandpass is small relative to $\lambda_r f/D$. In this case, the pixel intensity reduces to the simpler form,

$$I_{ij} \cong \eta^2 \Delta\lambda I_p(\lambda_r)\Delta\alpha\left[\frac{1}{1-(\Delta\lambda_r/2\lambda_r)^2}\right]P_{ij}, \quad (8)$$

[1] For *TPF-C* this widest band of interest, and that used for this ideal reference case, is planned to be from 0.5 to 0.8 μm. Strictly speaking, the variation of the field intensity, $I_p(\lambda)$, should be included in the integral over wavelength as the Planck curve has significant variations over this wide a band. For simplicity, however, we are approximating the response curve as flat here for the purposes of developing the metric. In the less idealized case in which we measure over narrow bands this approximation becomes much better.

where the unitless quantity,

$$P_{ij} = \frac{1}{\Delta\alpha}\int\int_{\Delta\alpha_{ij}} P\left(\lambda_r, u-\theta_x f, v-\theta_y f\right) du\, dv,$$

is the normalized integral of the PSF over pixel (i,j), and $\Delta\alpha$ is the pixel area. Since we are considering small wavelength bands, we can safely drop the term in brackets in equation (8).

The final electron count rate at each pixel (also Poisson distributed), including background, is then

$$\hat{I}_{ij} = \epsilon\eta^2\Delta\lambda I_p(\lambda_r)\Delta\alpha P(0,0)\bar{P}_{ij} + \epsilon\eta^2\Delta\lambda I_b\Delta\alpha, \qquad (9)$$

where ϵ is the quantum efficiency of the camera and we have normalized the PSF by its value at (0, 0) [so that $\bar{P}_{ij} = P_{ij}/P(0,0)$ and likewise $\bar{P}(\lambda,u,v) = P(\lambda,u,v)/P(\lambda,0,0)$]. We introduce the carat to differentiate between the photon and electron counts.

3.2. *The Contrast Ratio*

To facilitate later developments, we define here the contrast ratio, Q, first introduced by Brown & Burrows (1990),

$$Q \equiv \frac{I_p P(0,0)}{I_b}. \qquad (10)$$

In words, Q is the ratio of the peak of the planetary image (in photons s^{-1} m^{-2} μm^{-1}) to the local surface brightness of the background (in the same units).

Using the earlier expressions and considering only the suppressed starlight and exozodiacal background, the contrast ratio can be written

$$Q = \frac{1}{\mathcal{F}_{ez}\Omega_{ij}/I_p + \left(I_s/I_p\right)\left[P\left(\lambda,\theta_x f,\theta_y f\right)/P(0,0)\right]}. \qquad (11)$$

Letting Δmag be the difference in magnitudes between the star and planet, V the visual magnitude of the star, and ΔF the equivalent magnitude of the exozodi surface brightness, this expression simplifies to

$$Q = \frac{1}{10^{(V+\Delta\mathrm{mag}-\Delta F)/2.5}\Omega_{ij} + C10^{\Delta\mathrm{mag}/2.5}}, \qquad (12)$$

where C is the contrast between the peak of the PSF and the residual at the planet location (usually 10^{-10} by design). The local solar system zodiacal dust has a surface brightness of roughly 23 mag arcsec^{-2}. Most target stars are around 5 mag, and the minimum-brightness planet is assumed to have a Δmag of 25.

3.3. *Nondimensional Form*

It is often helpful, and very common, to nondimensionalize the image plane coordinates. Recalling that $P(\lambda_r,0,0) = A_c^2/\lambda_r^2 f^2$, the normalized PSF is given by

$$\bar{P}(u,v) = \frac{P(u,v)}{P(0,0)} = \frac{1}{A_c^2}\left|\int\int_S \mathcal{A}(x,y)e^{-j(2\pi/\lambda f)(xu+yv)}\, dx\, dy\right|^2. \qquad (13)$$

Next, introducing a reference length, D (the side of a square aperture or the diameter of a circular aperture, for instance), defining dimensionless pupil plane coordinates, $X = x/D$, $Y = y/D$, and a dimensionless angular position in the image plane in units of λ/D,

$$\xi = \frac{Du}{\lambda f}, \quad \zeta = \frac{Dv}{\lambda f},$$

gives the dimensionless form of the PSF,

$$\bar{P}(\xi,\zeta) = \frac{D^4}{A_c^2}\left|\int\int_{\bar{S}} \mathcal{A}(X,Y)e^{-j2\pi(X\xi+Y\zeta)}\, dX\, dY\right|^2. \qquad (14)$$

This equation can be modified slightly by introducing the definition $sD^2 \equiv A$. That is, the entrance aperture area can be related to the square of the reference length by the constant s. Thus, for example, for square apertures $s = 1$, for circular apertures $s = \pi/4$, and for elliptical apertures $s = \pi r/4$, where r is the aspect ratio of the ellipse. The normalized PSF in this notation is given by

$$\bar{P}(\xi,\zeta) = \frac{A^2}{A_c^2}\frac{1}{s^2}\left|\int\int_{\bar{S}} \mathcal{A}(X,Y)e^{-j2\pi(X\xi+Y\zeta)}\, dX\, dY\right|^2. \qquad (15)$$

Finally, (u,v) is also replaced in the expression for $\bar{P}_{ij}$ to get the pixel integrated PSF in dimensionless units,

$$\begin{aligned}\bar{P}_{ij} &= \frac{\lambda_r^2 f^2\Delta\bar{\alpha}}{\Delta\alpha D^2}\frac{1}{\Delta\bar{\alpha}}\int\int_{\Delta\bar{\alpha}_{ij}} \bar{P}(\xi,\zeta)\, d\xi\, d\zeta \\ &= \frac{1}{\Delta\bar{\alpha}}\int\int_{\Delta\bar{\alpha}_{ij}} \bar{P}(\xi,\zeta)\, d\xi\, d\zeta, \end{aligned} \qquad (16)$$

where $\Delta\bar{\alpha}$ is the dimensionless pixel area [$\Delta\bar{\alpha} = \left(D^2/\lambda_r^2 f^2\right)\Delta\alpha$]. Thus, as expected, $\bar{P}_{ij}$ can be computed in either dimensional or unitless coordinates.

Substituting for $P(0,0)$ and $\Delta\alpha$ into equation (9), the measurement equation can be rewritten

$$\hat{I}_{ij}(t) = \epsilon\eta^2\Delta\lambda I_p(\lambda_r)\Delta\bar{\alpha}\frac{A_c^2}{A}s\bar{P}_{ij} + \epsilon\eta^2\Delta\lambda I_b\Delta\alpha. \qquad (17)$$

Finally, we define the total throughput, $T = A_c/A$, the ratio of the exit pupil pseudoarea to entrance pupil area. Substituting into equation (17) leaves

$$\hat{I}_{ij}(t) = \epsilon\eta^2\Delta\lambda I_p(\lambda_r)\Delta\bar{\alpha}T^2 As\bar{P}_{ij} + \epsilon\eta^2\Delta\lambda I_b\Delta\alpha. \qquad (18)$$

Thus, the electron count has been rewritten in terms of the total throughput and dimensionless pixel area.

4. PLANET DETECTION

In this section we examine the detection problem via a linear combination of pixel measurements. The usual approach is to do PSF fitting, where the irradiance of the planet is estimated by fitting the PSF to the CCD image using standard photometric approaches (Brown 2005). In other words, a planet PSF fit is attempted at every pixel and the results compared to find a statistically significant outlier. A linear, unbiased procedure for PSF fitting is described in Burrows (2003) and summarized in the Appendix. Here we provide a probabilistic basis for using PSF

fitting for planet detection and develop unambiguous metrics. We also show that PSF fitting is, in fact, close to the optimal linear detection strategy.

4.1. *PSF Fitting*

As in Burrows (2003), we represent the Poisson counting process by a linear measurement model,

$$z_{ij} = C_p \bar{P}_{ij} + C_b + \nu, \tag{19}$$

where $C_p = \epsilon\eta^2 \Delta\lambda I_p(\lambda_r)\Delta\bar{\alpha}T^2 sAt$, $C_b = \epsilon\eta^2\Delta\lambda I_b \Delta\alpha t$, ν represents the photon and readout noise, and t is the integration time. Thus, C_p is approximately the mean photon count at the pixel centered on the PSF (since $P_{00} \sim 1$) and C_b is the mean background photon count at any single pixel. In addition, from equation (9) and the definition of the contrast ratio, Q, it is straightforward to show that $Q = C_p/C_b$.

Given this model in equation (19), a variety of methods are available to estimate C_p, the planet irradiance. The most common approach to photometry is via PSF fitting (also called "matched filtering"). The standard, linear method is described in Burrows (2003). We summarize this derivation in the Appendix and also compare it to maximum likelihood approaches.

While Burrows (2003) primarily considers the background-limited case, here we must assume that the planet signal and the background are of comparable intensity. (It is possible that the exozodi could be as much as 3–4 times the intensity of the planet. Nevertheless, this is not large enough to justify the background-limited regime, and thus it is best to include both the planet and background noises. Note that Brown [2005] uses the background-limited results of Burrows [2003] yet includes the planet signal, resulting in an inconsistency in his results.)

While it is possible to consider nonlinear maximum likelihood approaches to PSF fitting, these are complicated, provide minimal improvement (if any), are typically biased, and do not provide convenient expressions for the estimate and variance. We therefore consider only the linear, unbiased estimator suggested by Burrows (2003),

$$\hat{C}_p = \frac{\sum_{ij}(z_{ij} - C_b)\bar{P}_{ij}}{\sum_{ij}\bar{P}_{ij}^2}, \tag{20}$$

where the sum is taken over some region ΔS of the image plane covered by the PSF.

There is a critical assumption in this approach that the mean value of the background is known and constant across the pixels. This allows it to be subtracted in equation (20). While the uniformity assumption can be relaxed by indexing C_b for each pixel, its value still must be known. This is consistent with the assumptions in Brown (2005) and Burrows (2003), for example. One suggestion for avoiding the need to know the background is to take two images at two different rotation angles and subtract; this would eliminate the scattered and diffracted light background and the exozodi, as long as it is constant, at the expense of doubling the background photon noise. However, it does complicate the subsequent analyses, which assume that C_b is known for the signal-to-noise calculations.

In § 3 we explained that the background consisted primarily of the residual diffracted starlight, the scattered light from imperfections in the optics, and the exozodical light. Without some form of wavefront correction, the scattered light would dominate the background and overwhelm the planet by many orders of magnitude. Thus, we are assuming that the observatory has a wavefront control system that corrects the scattered light "speckles" down to the level of the other background sources. We also assume that the scattered light level is stable during an observation (a critical requirement for *TPF-C* as there is not an adequate signal to make corrections on faster timescales). Such an ability has been demonstrated in the laboratory.

These assumptions that the background is known and smooth and that the scattered light level has been corrected to a contrast at or below that of the diffracted light and zodiacal background, are strong ones. However, our goal was to produce the simplest meaningful analysis that could be compared to existing work and used to make baseline predictions regarding the performance and scope of a planet-finding mission. A more likely scenario is that the background is unknown and varies across the image plane and would thus have to be simultaneously estimated. In addition, the residual scattered light would most likely not be smooth and perhaps produce confusion in the PSF fitting process. Both of these situations can be treated via a more complicated, multidimensional estimation process using multiple rotations, different wavelengths, or both. We will address these issues in a future paper.

For pixels with a planet, equation (20) is an unbiased estimate of the planet signal, C_p, with variance

$$\sigma^2_{C_p} = \frac{C_p \sum_{ij}\bar{P}_{ij}^3}{\left(\sum_{ij}\bar{P}_{ij}^2\right)^2} + \frac{C_b}{\sum_{ij}\bar{P}_{ij}^2}. \tag{21}$$

For pixels without a planet, the expected value of equation (20) is zero and the variance is

$$\sigma_b^2 = \frac{C_b}{\sum_{ij}\bar{P}_{ij}^2}. \tag{22}$$

Planet detection is performed by calculating the estimate in equation (20) at each pixel in the dark hole region. For pixels without planets, the estimate will scatter about zero within the range defined by the variance in equation (22). For pixels with a planet, and a long enough integration time, the planet estimate should exceed this region about zero by some amount (see Fig. 4 for an illustrative example). We can plot the estimate for all pixels and look for where it exceeds the 1 σ bounds defined by equation (22) by some factor.

4.2. *Detection Metric and Missed Detection Probability*

Burrows (2003), and later Brown (2005), define the signal-to-noise ratio (S/N) as the ratio of the actual planet irradiance, C_p (the expected value of the estimate), and the standard deviation of the estimate error. While this is a reasonable indicator of the quality of the estimate, it says little about how to make a detection decision. The irradiance, C_p, is not the signal nor is the estimate variance the noise. This approach to computing integration time also ignores that fact that the estimate of C_p is also a random variable and thus has a specific, quantifiable probability of failing any detection test. Finally, Brown (2005) uses both the planet and background photon noise in the definition of the S/N. As just described, only the background is needed to perform a detection.

For the detection problem, we are interested in a metric that can be used to decide whether a planet exists in a given region. This can then be used to determine a maximum integration time before moving on. The signal in this case is the estimate of the irradiance, $\hat{C}_p$ (from eq. [20]), and the noise is given by the background standard deviation on pixels without a planet (from eq. [22]). Thus, a detection occurs when the estimate exceeds the background by

some amount. The resulting signal-to-noise ratio metric, $\hat{C}_p/\sigma_b$, is a random variable, and we can therefore define a confidence interval for the detection decision. We choose an integration time such that the probability that the S/N, assuming a planet exists, exceeds a given threshold, K, is equal to a desired value, P,

$$\Pr\left(\frac{\hat{C}_p}{\sigma_b} > K|\text{a planet}\right) = P. \tag{23}$$

For example, for 99% confidence, $P = 0.99$. That is, if the correct integration time is used and no planet is seen, then we can say with 99% confidence that there is no planet at that location (or at least that there is no planet within $K\sigma_b$ of the background noise).

To compute this probability, we first use equations (20) and (22) to find the S/N,

$$\frac{\hat{C}_p}{\sigma_b} = \frac{\sum_{ij}(z_{ij} - C_b)\bar{P}_{ij}}{\sqrt{C_b \sum_{ij} \bar{P}_{ij}^2}}. \tag{24}$$

Unfortunately, the probability mass function of $\hat{C}_p/\sigma_b$ is extremely difficult to compute using a Poisson process for the z_{ij}. However, for the integration times being considered, the arrival rate is high enough that the Poisson process is well approximated by a normal one, and this assumption allows for a straightforward, and adequate, estimate of the integration time. In addition, by the central limit theorem, the sum in equation (24) normalizes the random variable. From equation (24) the mean and variance are easily computed,

$$\mu_{\mathrm{S/N}} = \sqrt{C_p Q \sum \bar{P}_{ij}^2}, \tag{25}$$

$$\sigma_{\mathrm{S/N}}^2 = 1 + \frac{Q \sum \bar{P}_{ij}^3}{\sum \bar{P}_{ij}^2}, \tag{26}$$

where we have used the fact that $Q = C_p/C_b$.

This allows the probability in equation (23) to be rewritten in terms of a normalized variable, Z,

$$\Pr\left(Z > \frac{K - \mu_{\mathrm{S/N}}}{\sigma_{\mathrm{S/N}}}|\text{a planet}\right) = P, \tag{27}$$

where Z is a standard normal random variable with zero mean and unit variance. The detection criterion is then given by the probability integral,

$$1 - \Phi\left(\frac{K - \mu_{\mathrm{S/N}}}{\sigma_{\mathrm{S/N}}}\right) = P \tag{28}$$

or

$$\Phi\left(\frac{K - \mu_{\mathrm{S/N}}}{\sigma_{\mathrm{S/N}}}\right) = 1 - P = P_{\mathrm{MD}}, \tag{29}$$

where $P_{\mathrm{MD}} = 1 - P$ is the probability of a missed detection and $\Phi(z)$ is the Gaussian distribution function,

$$\Phi(z) = \Pr(Z \le z) = \int_{-\infty}^{z} \frac{1}{\sqrt{2\pi}} e^{-y^2/2}\, dy = \frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{z}{\sqrt{2}}\right)\right]. \tag{30}$$

In other words, the integration time is chosen so that $(K - \mu_{\mathrm{S/N}})/\sigma_{\mathrm{S/N}}$ reaches a value, γ, such that $\Phi(\gamma) = 1 - P$. For example, a 0.1% missed detection rate ($P_{\mathrm{MD}} = 0.001$) corresponds to $\gamma = -3.1$. The detection criterion can now be written

$$\frac{K - \sqrt{C_p Q \sum \bar{P}_{ij}^2}}{\sqrt{1 + \left(Q \sum \bar{P}_{ij}^3\right)\left(\sum \bar{P}_{ij}^2\right)^{-1}}} = \gamma. \tag{31}$$

To solve for the needed integration time to achieve the desired false alarm and missed detection rates, this equation is first solved for C_p (since Q is independent of t) and then divided through to find the integration time,

$$t = \frac{1}{\epsilon\eta^2 \Delta\lambda I_p(\lambda_r) T^2 A s} \frac{\left[K - \gamma\sqrt{1 + \left(Q \sum \bar{P}_{ij}^3\right)\left(\sum \bar{P}_{ij}^2\right)^{-1}}\right]^2}{Q \Delta\bar{\alpha} \sum \bar{P}_{ij}^2}. \tag{32}$$

For consistency with Brown & Burrows (1990) and Burrows (2003), and for computational convenience, we make one more modification by introducing the "sharpness." Sharpness is defined in Burrows (2003) as $\Psi = \sum_{ij} \bar{P}_{ij}^2/(\sum \bar{P}_{ij})^2$. Using sharpness desensitizes the calculation to the location of the PSF within a pixel. Sharpness is a unique characteristic of the PSF of any coronagraph. With this substitution, equation (32) becomes

$$t = \frac{1}{\beta} \frac{\left(K - \gamma\sqrt{1 + \tilde{Q}\Xi/\Psi}\right)^2}{\tilde{Q} T_A \Psi}, \tag{33}$$

where we have also defined $\beta = \epsilon\eta^2 \Delta\lambda I_p(\lambda_r) TA$, $\tilde{Q} \equiv Q \sum_{ij} \bar{P}_{ij}$, $\Xi \equiv \sum \bar{P}_{ij}^3 (\sum \bar{P}_{ij})^{-3}$, and

$$T_A \equiv Ts\Delta\bar{\alpha} \sum \bar{P}_{ij} = Ts \int\int_{\Delta S} \bar{P}(\xi, \zeta)\, d\xi\, d\zeta.$$

The parameter $\tilde{Q}$ is the ratio of the total photon count from the planet over the integration area, ΔS, to the number of background counts in a single pixel, and T_A is what we called the "Airy throughput" in Vanderbei et al. (2003b). It is the throughput times the integral of the normalized PSF over the integration area ΔS. It is straightforward to show, using Parseval's theorem, that $s \int\int_{\Delta S} \bar{P}(\xi, \zeta)\, d\xi\, d\zeta = 1$ for any open aperture assuming that ΔS covers the entire image plane. For the expected smaller integration area, T_A is always less than T. It is thus the fraction of the total throughput achieved for the given region of the PSF used in detection. This difference can be significant. For an open, circular aperture, the Airy throughput corresponding to the first null is only 84.5%. If only the FWHM is used, then this drops to 53%, almost a factor of 2 in integration time. For shaped pupils, this is also important, since most shaped pupil PSFs move a significant portion of the light outside the "outer working angle"; this is light not used in the planet detection.

This dependence on the Airy throughput would imply that it is always best to use as many pixels as possible to cover as much of the planet's PSF as possible. However, the remaining terms that depend on $\tilde{Q}$, Ψ, and Ξ have the opposite trend, suggesting that a smaller number of pixels around the central core is better (since more pixels tend to include more background light but less

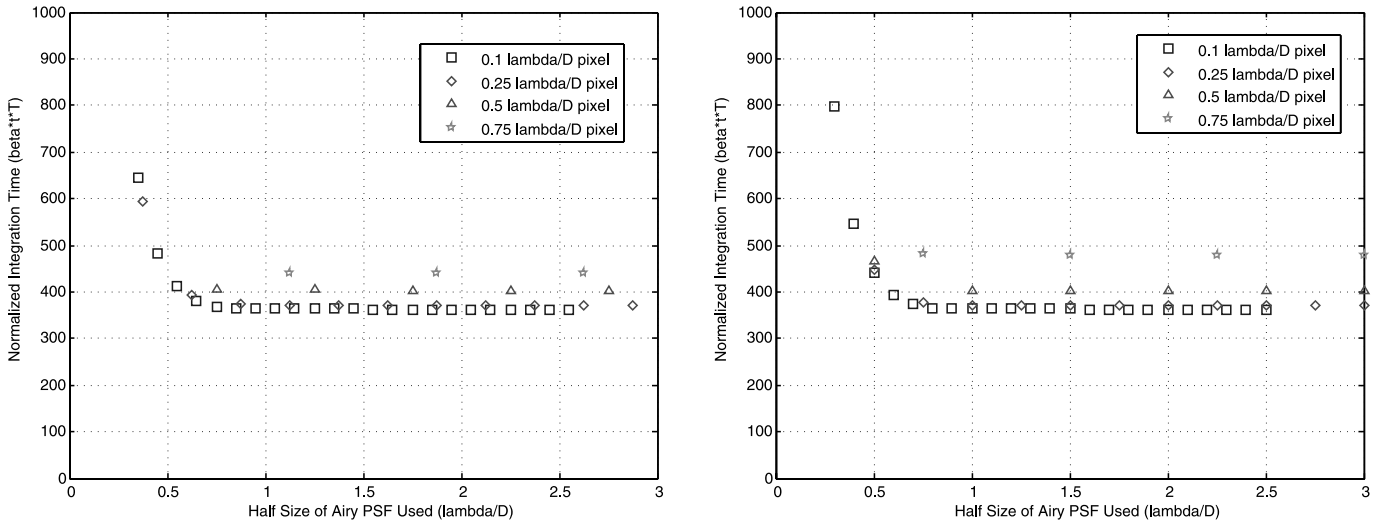


FIG. 1.—Normalized integration time vs. size of Airy PSF used in detection for various pixel sizes. *Left*: PSF centered on single pixel. *Right*: PSF centered on corner of four pixels.

planet light because of the drop-off of the PSF). The result is that beyond a certain region of the core there is no improvement in integration time. This region size depends on the specific PSF and the size of the pixels. For an Airy PSF it is roughly $0.75\lambda/D$. Figure 1 shows two plots of the normalized integration time (βt) versus width of the core used in the detection for four different pixel sizes and a background-limited detection, where $Q = \frac{1}{3}$. While there is an advantage to oversampling the PSF, the gain is slight. For smaller background, and larger Q, slightly less of the core is necessary but the difference in integration time is also small. We should note that this assumes a fixed Q, that is, that the background is independent of pixel size. In fact, because the exozodi level depends on Ω_{ij}, for large exozodis a smaller pixel size (and oversampled PSF) can dramatically reduce the background, increase Q, and thus reduce the integration time.

This also raises the question of the location of the PSF within the pixel. Up to now, this has been ignored. Figure 1 shows that by using sharpness the integration time expression is reasonably insensitive to the PSF location relative to the pixel center. However, in practice, the estimate of C_p in equation (20) assumes that P_{ij} is known at every pixel, introducing an ambiguity. One solution is to attempt a fit for many values of P_{ij} as the PSF moves relative to a pixel center. Alternatively, we could compute the sensitivity of the detection to errors in P_{ij} and increase the integration time accordingly. Finally, we could treat the actual value of P_{ij} as a random variable with uniform distribution between its value for the PSF landing on a pixel center and its value for the PSF landing on a pixel edge. This probability distribution can then be included in the probability calculation in equation (23). We defer this approach to future work.

4.3. *Choice of K and False-Alarm Probability*

Equation (33) provides the integration time to achieve a desired small missed detection rate for a given detection threshold, K. However, the choice of K still seems arbitrary. Most astronomical applications select K between 5 and 10 (while ignoring the missed detection probability embedded in γ). In reality, K has a probabilistic basis as well; it is a measure of the probability of false alarms. In fact, the best approach is to pick the integration time, t, and the threshold, K, to achieve a desired balance between the two types of errors.

The probability of false alarms is given by an expression similar to that in equation (23),

$$\Pr\left(\frac{\hat{C}_p}{\sigma_b} > K|\text{no planet}\right) = P_{\text{FA}}. \tag{34}$$

The expression for $\hat{C}_p/\sigma_b$ is the same as that in equation (24), but since there is no planet the mean and variance differ,

$$\mu_{\text{FA}} = 0, \tag{35}$$

$$\sigma^2_{\text{FA}} = 1. \tag{36}$$

The S/N when there is no planet in a pixel is, in fact, well approximated by a standard normal random variable. The false-alarm probability is then given simply by

$$\Phi(K) = 1 - P_{\text{FA}}. \tag{37}$$

Thus, equation (37) is used to select the threshold, K, corresponding to a desired probability of false alarms, and equation (33) is then used to select the integration time to achieve a desired probability of missed detections.

Since we expect at best to only find one planet per system surveyed, most pixels will not have a planet. This places a high premium on keeping the false-alarm rate low (a large K), otherwise much time would be wasted on follow-up of false planets. Nevertheless, a K-value greater than 5 results in a false-alarm probability lower than necessary [e.g., $\Phi(5) = 0.9999997$, a false-alarm probability of 3×10^{-7}]. The result is a potentially longer integration time than necessary, dramatically affecting the mission. A reasonable approach is to select the false-alarm probability such that there is less than one false alarm per observed system. For instance, for a *TPF* mission that surveys 100 systems, a total false-alarm rate of 1% would result in one false alarm over the course of the mission. Since each planet detection will be followed by deep surveys for characterization, the false alarm would be quickly discovered and have minimal impact on mission life. If we assume on the order of 3000 measurements in a given system, we therefore require that $3000P_{\text{FA}} \sim 0.1$ or $P_{\text{FA}} \sim 3 \times 10^{-5}$. This corresponds to $K = 4$.

The missed detection probability is critical for a mission with so few surveyed systems and unknown planet occurrence rate, $\eta_\oplus$. It should be made low enough to avoid a single missed detection over the course of the mission. A missed detection probability of less than 0.1% is most likely needed, requiring $\gamma = -3.1$.

With these numbers, the integration time equation becomes

$$t = \frac{1}{\beta}\frac{\left(4 + 3.1\sqrt{1 + \tilde{Q}\Xi/\Psi}\right)^2}{\tilde{Q}T_A\Psi}. \tag{38}$$

There are a number of notable differences between this expression for integration time and that in Brown (2005). First, Brown (2005) selected the S/N threshold by implicitly only considering the false-alarm probability. He ignored the missed detection probability (γ) and the potentially significant contribution to the integration time in equation (33). (Physically, this corresponds to basing the detection decision only on the mean value of the S/N that leaves a significant probability that a particular S/N may fall below the threshold. Mathematically, this means $\gamma = 0$, which corresponds to a missed detection probability of 50%.) Brown (2005) also neglected to account for the possible reduction in Airy throughput due to only using the center of the PSF.

Another difference is that Brown (2005) included the planet photon noise in the S/N expression. We describe above why only the background noise should be used in the S/N definition.

Also note that the needed integration time depends on the actual planet irradiance (in β). Strictly speaking, this is itself a random variable as the planet irradiance could vary quite widely among systems. While it is possible to include a density function for I_p and marginalize the S/N, this makes the problem substantially more complicated. It is also not clear that an a priori density for the planet irradiance can easily be found. A more reasonable approach is to use a worst-case irradiance in equation (33) for the most conservative integration time, at the possible cost of an unnecessary increase in mission length. This will become important again in the Bayesian analysis later.

4.4. *Numerical Example*

In this section we present a simple numerical calculation of the integration time for a typical system. For simplicity, we ignore the ellipticity of the *TPF-C* entrance aperture. To begin, we also ignore the coronagraph (and any affect on the PSF) and consider the shortest possible integration time to detect a planet given the telescope aperture. In other words, $A_c = A$, the entrance aperture area. We assume the entrance area, A, is the same as the 8 m × 3.5 m *TPF-C* design (22 m^2). Assuming a 5 mag star and a planet with $\Delta\text{mag} = 25$, the irradiance is approximately $I_p = 9.5 \times 10^{-9}$ photons cm^{-2} nm^{-1} s^{-1}. With a quantum efficiency of 0.8, $\Delta\lambda = 100$ nm, and $\eta^2 = 0.33$, $\beta = 0.06$ photons s$^-$. The integration time then becomes

$$t \cong 18\frac{\left(4 + 3.1\sqrt{1 + \tilde{Q}\Xi/\Psi}\right)^2}{\tilde{Q}T_A\Psi}. \tag{39}$$

Assuming a fairly conservative value of $Q = \frac{1}{3}$ (large exozodi), we can use Figure 1 to obtain the normalized integration of roughly 400 for a critically sampled Airy function. The resulting integration time becomes 7200 s or 2 hr.

A more realistic case would include the effect of a Lyot stop (or shaped pupil). The Lyot stop reduces T and T_A (by 0.3 for the current design). If we assume that a detector with larger pixels is used so that the Airy function is still critically sampled, the sharpness, Ψ, Ξ, and $\tilde{Q}$ are all unchanged. The result is an increase in the integration time by a factor of $1/0.3^2$ or about 11. This implies a required integration time of 22 hr.

This is only a single integration example assuming a planet is visible in the discovery space of the telescope. In other words, an observatory with a coronagraph designed to achieve a contrast of 10^{-10}, looking at a 5 mag star 10 pc away with a 1 AU terrestrial planet at maximum elongation, will have a 99.9% chance of discovering that planet in less than 22 hr of observing. These calculations are critical for any analysis to plan a mission and study the completeness of coverage, including the ephemerides of each planet and the probability they are within the discovery space. Such an analysis, however, is outside the scope of this paper. This is discussed in part in Brown (2005). In this paper we have provided the tools to update and expand this treatment and place it on a sound statistical footing.

4.5. *General Linear Detection and Optimality*

The specific irradiance estimate in equation (20) can be easily generalized into a form that covers all possible linear processing techniques,

$$Z = \sum_{ij} W_{ij}\left(z_{ij} - C_b\right), \tag{40}$$

for some unknown set of weights, W_{ij}, over a region ΔS of the image. Following the same process as before, the signal-to-noise ratio is

$$\frac{Z}{\sigma_b} = \frac{\sum_{ij} W_{ij}\left(z_{ij} - C_b\right)}{\sqrt{C_b \sum_{ij} W_{ij}^2}}, \tag{41}$$

with mean $\mu_{\text{S/N}} = (C_p Q)^{1/2} \sum_{ij} W_{ij}\bar{P}_{ij}/(\sum_{ij} W_{ij}^2)^{1/2}$ and variance $\sigma^2_{\text{S/N}} = 1 + (Q\sum W_{ij}^2\bar{P}_{ij})(\sum W_{ij}^2)^{-1}$. By the same probabilistic argument, the detection condition becomes

$$\frac{K - \left(\sqrt{C_p Q}\sum W_{ij}\bar{P}_{ij}\right)\left(\sqrt{\sum W_{ij}^2}\right)^{-1}}{\sqrt{1 + \left(Q\sum W_{ij}^2\bar{P}_{ij}\right)\left(\sum W_{ij}^2\right)^{-1}}} = \gamma. \tag{42}$$

As before, this can be solved for C_P and then t,

$$t = \frac{1}{\beta}\frac{\left[K - \gamma\sqrt{1 + \left(Q\sum W_{ij}^2\bar{P}_{ij}\right)\left(\sum W_{ij}^2\right)^{-1}}\right]^2 \sum W_{ij}^2}{Q\Delta\bar{\alpha}\left(\sum W_{ij}\bar{P}_{ij}\right)^2}. \tag{43}$$

One observation from equation (43) is that the weights can be multiplied by an arbitrary constant without effecting the integration time (and thus the detection). This is because any constant weight divides out when forming the S/N statistic. Thus, while the constant term $\sum_{ij}\bar{P}_{ij}^2$ in equation (20) is important for an unbiased estimate of C_p (photometry), it is not necessary for a simple detection.

Equation (43) can be used to find the optimal weights resulting in a minimum integration time. These turn out to be very close to PSF fitting in § 4.1.

This minimum is found by differentiating equation (43) with respect to W_{ij} and setting equal to zero. Unfortunately, the exact problem is difficult to solve because of the term in the radical. However, for $Q \ll 1$, this term is small and the integration time is very well approximated by

$$t \cong \frac{1}{\beta}\frac{(K-\gamma)^2 \sum W_{ij}^2}{Q\Delta\bar{\alpha}\left(\sum W_{ij}\bar{P}_{ij}\right)^2}. \tag{44}$$

The resulting minimum is found from

$$\frac{\partial}{\partial W_{ij}}\left[\frac{\sum W_{ij}^2}{\left(\sum W_{ij}\bar{P}_{ij}\right)^2}\right] = 0. \tag{45}$$

Completing the differentiation results in the condition for W_{ij},

$$W_{ij} = P_{ij}\frac{\sum W_{ij}^2}{\sum W_{ij}\bar{P}_{ij}}. \tag{46}$$

While difficult to solve in general, it is easy to show that this equation is satisfied by the PSF fitting weights used in § 4.1, $W_{ij} = \bar{P}_{ij}$. Thus, for small Q (large background), the PSF fitting approach is the approximately optimal linear detection technique. Any improvement from searching for the exact optimal weights is only very slight.

5. BAYESIAN DETECTION

We showed above that PSF fitting for detection was close to the optimal, that is, shortest time, detection among all linear algorithms. It can also be shown that under certain assumptions it achieves the Cramer-Rao bound and is thus the best possible estimate of the irradiance. However, for the detection problem we are not necessarily interested in planetary photometry, just a determination of whether a planet exists. PSF fitting may then provide more information than is necessary. It is reasonable, therefore, to ask whether there might be faster algorithms for making the detection decision. In this section we explore an alternative hypothesis-testing approach using Bayesian methods with the same measurement model as in § 3.

5.1. *Detection Criterion*

To determine whether a potential planet is present in a set of pixels ΔS, centered at pixel n, we compare the probabilities of the following two alternate hypotheses:

$$H_0^n: \text{ there is no planet in the } n\text{th pixel},$$
$$H_1^n: \text{ there is a planet in the } n\text{th pixel}.$$

Each hypothesis, H_i, has the associated model M_i:

$$M_0^n: \ \lambda_{ij} = C_b,$$
$$M_1^n: \ \lambda_{ij} = C_p\bar{P}_{ij} + C_b,$$

where λ_{ij} is the mean arrival rate of photons in the Poisson distribution at pixel (i,j) located in ΔS.

We now compare the posterior probabilities of the two models conditioned on the collected data and decide that there is a planet in pixel n if the "odds ratio," O_{10}^n, favors H_1^n over H_0^n:

$$O_{10}^n = \frac{p\left(M_1^n|Z\right)}{p\left(M_0^n|Z\right)} = \frac{p\left(M_1^n\right)p\left(Z|M_1^n\right)}{p\left(M_0^n\right)p\left(Z|M_0^n\right)} = R\frac{p\left(Z|M_1^n\right)}{p\left(Z|M_0^n\right)}, \tag{47}$$

where Z is the set of collected measurements and we have used Bayes's rule to write the odds ratio as the product of the ratio of prior probabilities, $R = p\left(M_1^n\right)/p(M_0^n)$, and the ratio of likelihoods, $p\left(Z|M_1^n\right)/p\left(Z|M_0^n\right)$. The dependence on the ratio of prior probabilities of there being a planet is troubling at first, since it appears that the effectiveness of the test can be greatly affected by this essentially unknown quantity (alternatively, one can view this as perhaps an attractive feature as prior information on the frequency becomes available). We show below that this dependence is easily removed by balancing the sensitivity between missed detections and false alarms.

One possible and simple decision criterion would be that $O_{10}^n > 1$ favors H_1^n, thus indicating that there is a planet whose true location ζ^* is inside ΔS. However, a better approach to the decision is to base it on a different value of the ratio derived from particular weighting criteria. This is called a "loss-based decision." The specific value of the odds ratio for the decision is selected to balance the relative priorities between missed detections and false alarms (note the similarity to the PSF fitting test above). In most applications this process is used to optimize the power of the test (via the receiver operator curve [ROC]). Here, however, there is a large incentive to reduce the probability of false alarms well below that of missed detections rather than balance the two types of errors. In what follows, we develop analytical expressions to quantitatively trade off these two probabilities.

5.2. *Likelihoods and Priors*

In general, the probability of a given data set depends on various physical parameters of the system. For instance, the mean photon rate in the Poisson distribution depends on the irradiance of the planet and its location in the pixel. As mentioned in § 4, this is, strictly speaking, a random variable. The usual approach to computing the likelihood functions is to *marginalize* the probabilities to remove the parameter dependence altogether. However, given our lack of knowledge regarding the planet distribution, and that our goal herein is to simply find a conservative estimate of the needed integration time for detection, it is more sensible to begin with a worst-case analysis. Thus, we assume the parameters (such as planet irradiance) are known and take on their most pessimistic value (for instance, for planets below a certain limiting Δmag we assume a detection is impossible). If an actual planet turns out to be brighter, then we have simply integrated longer than necessary, but we will not miss the detection. In short, we trade conservatism for our lack of knowledge of the specific distribution of the priors.

The odds ratio now becomes the ratio of the likelihood functions times the prior probabilities,

$$O_{10}^n = R\frac{L_1}{L_0}, \tag{48}$$

where the likelihood functions are given by $L_1 = p(Z|M_1^n)$ and $L_0 = p(Z|M_0^n)$. Each of these probability distributions per pixel are Poisson with parameter λ_{ij},

$$L_l = \prod_{ij}\frac{1}{z_{ij}!}\lambda_{ij}^{z_{ij}}e^{-\lambda_{ij}}. \tag{49}$$

Thus, for a set of data per pixel, z_{ij}, we form the odds ratio using our worst-case value for the planet irradiance, $\bar{C}_p$, and the contrast ratio, $\bar{Q}$, where we have used the bar to contrast the

value used in the detection algorithm from the true values, C_p and Q. The background irradiance, C_b, is still assumed to be known (and uniform). Substituting into equation (48), the odds ratio becomes

$$O_{10}^n = Re^{-\bar{C}_p \Delta P} \prod_{ij} \left(1 + \bar{Q}\bar{P}_{ij}\right)^{z_{ij}}, \tag{50}$$

where $\Delta P = \sum_{ij} \bar{P}_{ij}$. It is often more convenient to use the log of the odds ratio as the test statistic,

$$\log O_{10}^n = \log R - \bar{C}_p \Delta P + \sum_{ij} z_{ij} \log\left(1 + \bar{Q}\bar{P}_{ij}\right). \tag{51}$$

The two constant terms do not affect the statistical test, so it is easier to define a new test statistic that removes them,

$$\chi^n = \sum_{ij} z_{ij} \log\left(1 + \bar{Q}\bar{P}_{ij}\right). \tag{52}$$

Unlike in PSF fitting, not only do we need to know P_{ij}, we also need to know Q. This forces us to ask about the sensitivity of the method to errors in our assumed knowledge, $\bar{Q}$. As mentioned above, for now we use a worst-case approach, using the smallest value of I_p expected (recall that in both of these approaches the background is assumed known).

The test statistic, χ, depends on the measured data and assumed, worst-case parameters. What remains is to find a threshold value for χ to decide between the two hypotheses.

5.3. *Hypothesis Tests*

Again, the loss-based decision is made based on the value of the test statistic. That is, just as before, if $\chi^n > K$ then we reject H_0 and decide that a planet is present in ΔS. Otherwise, we accept H_0 and assume no planet.

5.3.1. *Missed Detections*

The probability of a missed detection can be written formally as

$$\Pr\left(\chi^n < K | M_1^n\right) = P_{\rm MD}$$
$$\text{or}$$
$$\Pr\left(\sum z_{ij} \log\left(1 + \bar{Q}\bar{P}_{ij}\right) < K | M_1^n\right) = P_{\rm MD}. \tag{53}$$

To compare to PSF fitting, we look for an expression that allows us to pick an integration time to achieve a certain desired probability of missed detection, $P_{\rm MD}$. Unfortunately, even knowing that z_k is Poisson, forming the probability distribution in equation (53) is formidable. For a precise result, we must resort to Monte Carlo simulations. However, for the expected integration times and arrival rates here, the Poisson distribution is close to Gaussian, and as in § 4, this assumption is useful for obtaining integration time estimates. Thus, letting $B_{ij} = \log\left(1 + \bar{Q}\bar{P}_{ij}\right)$, the probability in equation (53) can be rewritten in terms of a standard normal random variable,

$$\Pr\left(Z < \frac{K - \mu_Z}{\sigma_Z}\right) = P_{\rm MD}, \tag{54}$$

where $\mu_Z = C_p \sum B_{ij}(\bar{P}_{ij} + 1/Q)$, $\sigma_Z^2 = C_p \sum B_{ij}^2(\bar{P}_{ij} + 1/Q)$, and Z is normally distributed with mean 0 and standard deviation 1. Thus, the missed detection equation becomes

$$\Phi\left(\frac{K - \mu_Z}{\sigma_Z}\right) = P_{\rm MD}, \tag{55}$$

where, as with PSF fitting, the desired missed detection probability, $P_{\rm MD}$, is picked a priori. Letting γ again be the value of the argument such that $\Phi(\gamma) = P_{\rm MD}$, the integration time criterion becomes

$$\frac{K - \mu_Z}{\sigma_Z} = \gamma. \tag{56}$$

In other words, for a desired missed detection probability, the integration time is selected such that equation (56) is satisfied.

Substituting for μ_Z and σ_Z leaves

$$K - C_p \sum B_{ij}\left(\bar{P}_{ij} + \frac{1}{Q}\right) = \gamma \sqrt{C_p \sum B_{ij}^2\left(\bar{P}_{ij} + \frac{1}{Q}\right)}. \tag{57}$$

This equation has two unknowns, K and t. As before, the false-alarm rate provides a criterion for the threshold, K, that can be substituted back in to find the integration time.

5.3.2. *False Alarms*

The probability of a false alarm is similar to equation (53),

$$\Pr\left(\chi^n > K | M_0^n\right) = P_{\rm FA}$$
$$\text{or}$$
$$\Pr\left(\sum z_{ij} \log\left(1 + \bar{Q}\bar{P}_{ij}\right) > K | M_0^n\right) = P_{\rm FA}. \tag{58}$$

As before, this is rewritten in terms of a standard normal random variable,

$$\Pr\left(Z' > \frac{K - \mu_{Z'}}{\sigma_{Z'}}\right) = 1 - \Phi\left(\frac{K - \mu_{Z'}}{\sigma_{Z'}}\right) = P_{\rm FA}, \tag{59}$$

where here $\mu_{Z'} = \sum B_{ij} C_b$ and $\sigma_{Z'}^2 = \sum B_{ij}^2 C_b$. Letting δ be the value of the argument such that $\Phi(\delta) = 1 - P_{\rm FA}$ leaves

$$\frac{K - \mu_{Z'}}{\sigma_{Z'}} = \delta \tag{60}$$

or

$$K = \delta \sqrt{\sum B_{ij}^2 C_b} + \sum B_{ij} C_b. \tag{61}$$

This equation provides the detection threshold test for χ^n for a given false-alarm probability (embodied in the value for δ).

5.3.3. *Integration Time*

The Bayesian integration time can now be found by using the false-alarm–based expression for K in equation (61) in equation (57),

$$C_p \left[\frac{\sum B_{ij}}{Q} - \sum B_{ij}\left(\bar{P}_{ij} + \frac{1}{Q}\right)\right] = \gamma \sqrt{C_p \sum B_{ij}^2\left(\bar{P}_{ij} + \frac{1}{Q}\right)} - \delta \sqrt{C_p \frac{\sum B_{ij}^2}{Q}}. \tag{62}$$

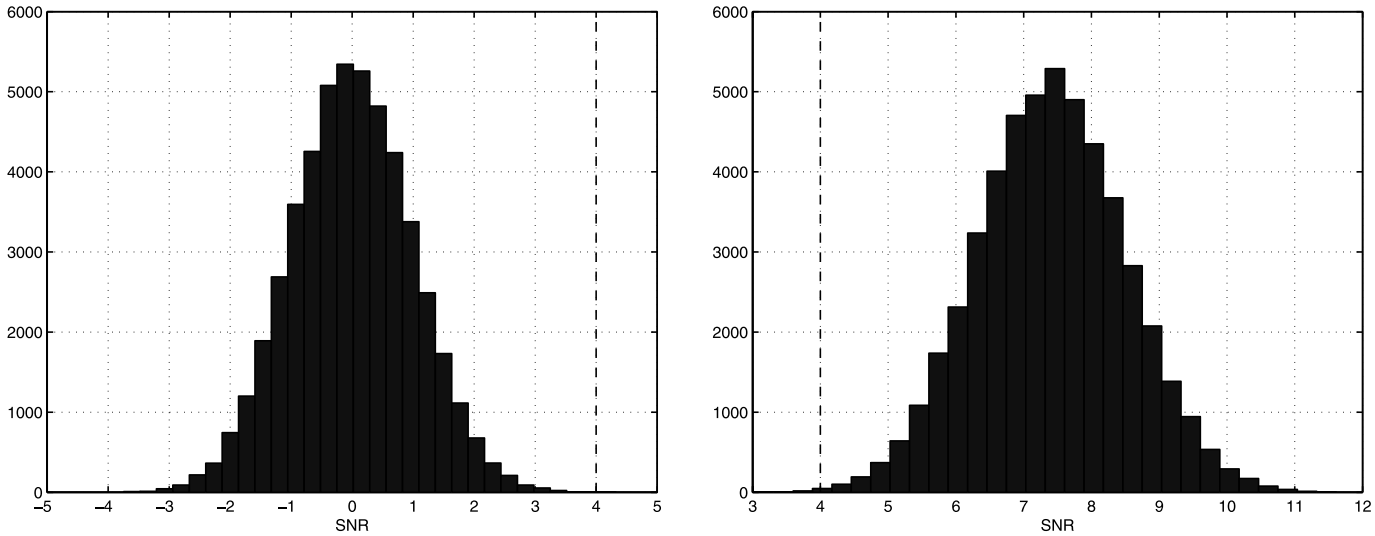


FIG. 2.—Histograms of 50,000 Poisson simulations of the PSF fitting signal-to-noise ratio test statistic for each case, pixels without a planet (*left*) and pixels with a planet (*right*). Dashed line indicates the threshold test value, K.

This can be used to solve for C_p as before and subsequently the Bayesian integration time,

$$t_{\mathrm{B}} = \frac{1}{\epsilon \Delta\lambda I_p(\lambda_r) T A_c s} \frac{\left[\gamma\sqrt{\sum B_{ij}^2\left(Q\bar{P}_{ij}+1\right)} - \delta\sqrt{\sum B_{ij}^2}\right]^2}{Q\Delta\bar{\alpha}\left(\sum B_{ij}\bar{P}_{ij}\right)^2}. \tag{63}$$

Finally, we can use the same renormalization as before and write the Bayesian integration time in terms of β, T_A, and $\tilde{Q}$,

$$t_{\mathrm{B}} = \frac{1}{\beta} \frac{\left[\gamma\sqrt{\sum B_{ij}^2\left(\tilde{Q}\bar{P}_{ij}/\sum\bar{P}_{ij}+1\right)} - \delta\sqrt{\sum B_{ij}^2}\right]^2}{\tilde{Q}T_A\left(\sum B_{ij}\bar{P}_{ij}/\sum\bar{P}_{ij}\right)^2}. \tag{64}$$

Surprisingly, for background-limited systems in which Q is small, there is no measurable improvement in integration time with the Bayesian approach. This is rather easy to show using equation (64). For small Q, $B_{ij} = \log(1 + \bar{Q}P_{ij}) \cong \bar{Q}P_{ij}$. With this substitution, equation (64) reduces exactly to equation (33). Thus, this Bayesian approach, in which we assumed we knew the background and used a worst-case value of I_p (i.e., Q), provides no integration time benefit over PSF fitting for small Q (large background). In addition, one would expect a poorer performance (that is, a slight increase in the probability of false alarms and missed detections) when $\bar{Q} \neq Q$. In addition, this approach suffers from the same ambiguity regarding the location of the PSF relative to a pixel center as PSF fitting, as P_{ij} has been assumed known throughout.

On the other hand, one of the attractive features of the Bayesian approach is that it provides a methodology for handling these unknowns. For example, the uncertainty in $\bar{Q}$ can be removed by assuming a distribution for I_p and marginalizing the likelihood function. Likewise, the ambiguity with regard to pixel size can also be removed via marginalization. Finally, we can study the case in which the background is not known by again marginalizing

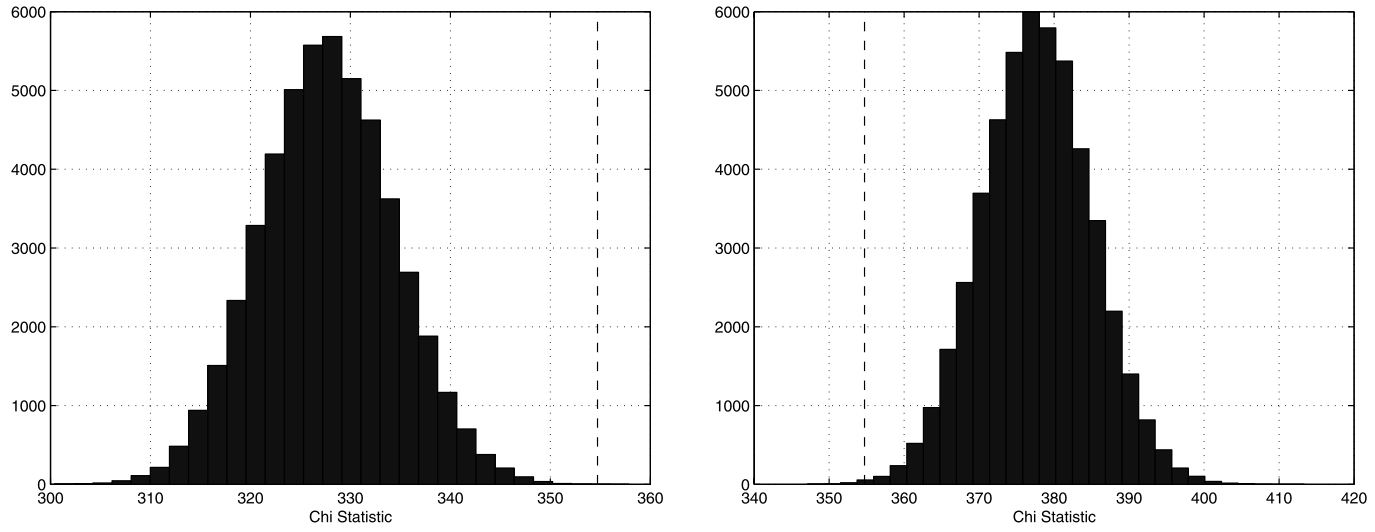


FIG. 3.—Histograms of 50,000 Poisson simulations of the Bayesian test statistic for each case, pixels without a planet (*left*) and pixels with a planet (*right*). Dashed line indicates the threshold test value, γ.

over a distribution for the background photon count. An unknown background can also be treated in PSF fitting by simultaneously estimating the planet irradiance and background, at the cost of much greater complexity and much longer integration times. We defer these comparisons to a future paper.

6. MONTE CARLO SIMULATION

To verify these results, we ran 50,000 Monte Carlo simulations for each case, with a planet and without. The Poisson arrival rate was computed using the integration time values computed in equations (38) and (64). An Airy PSF was used for all simulations, and Q was set to $\frac{1}{3}$. Figure 2 shows histograms of the results of the signal-to-noise test statistic using PSF fitting, and Figure 3 shows the equivalent results for the Bayesian test statistic, χ^n. Also shown on the plots are dashed lines indicating the threshold test levels for the false-alarm and missed detection probabilities described above. In both cases, there were 40 missed detections for the 50,000 runs with planets. For PSF fitting there were three false alarms, and in the Bayesian approach there were two. These plots confirm the analytical results and also justify the Gaussian assumption.

Figure 4 shows 200 samples each of the PSF fitting test statistic with and without a planet. This illustrates how the test statistic indicates the presence of a planet by exceeding the zero mean scatter by a threshold amount.

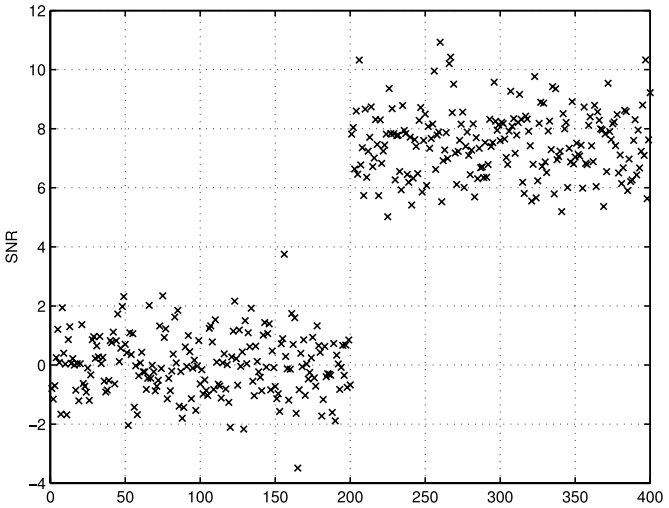


FIG. 4.—Samples of the simulated PSF fitting test statistic with and without a planet for the selected integration time.

7. FINAL REMARKS

In this paper we developed a careful model of the coronagraph measurement and used it to define two different detection approaches, matched filtering and Bayesian. For both approaches, we developed a simple expression for the needed integration time to achieve the desired probability of success. These integration time results can now be used for mission completeness studies.

The story is far from complete, however. Critical to these studies is an understanding of the expected background and the best possible estimates of Q. In addition, we have ignored the question of wavefront error and control. In this simplified model, the background was assumed known and uniform. In reality, the background is going to be a mix of exozodiacal light and residual wavefront error. The speckle background due to uncorrected wavefront error will not be uniform and can potentially confuse the planet detection algorithm. While the planet and background can be simultaneously estimated, it complicates the algorithm and lengthens the integration time. In addition, the additional flexibility of multiwavelength measurements has not been explored. Finally, the sensitivity questions of mismatches in assumed parameters and PSF alignment on pixels has not yet been studied. These issues are all currently being addressed and will be presented in a future paper. Nevertheless, the fundamental model and technique presented here provides an important benchmark to measure the potential search capabilities of *TPF*.

We gratefully acknowledge the support of the National Aeronautics and Space Administration through the Jet Propulsion Laboratory, California Institute of Technology for this work. We also acknowledge the support of I. Braems under a fellowship from the Institut National de Recherche en Informatique et Automatique (INRIA).

APPENDIX

PHOTOMETRY VIA PSF FITTING

In this appendix we briefly summarize the derivation of equation (20). This is a slightly expanded treatment of the same problem as in Burrows (2003). It also demonstrates the overdetermined, maximum likelihood estimation (MLE) case in Lyon et al. (1997). The problem is to estimate the irradiance, or photon count, of a signal corresponding to a point source in a known, uniform background. The image is given by the measurement $I_{ij} = B + AP_{ij} + N_{ij}$, where B is the assumed known background, A is the source irradiance, P_{ij} is the point-spread function spread over the pixels (i,j), and N_{ij} are noise photons in each pixel (i,j) with zero mean and variance σ_{ij}^2. The noise is assumed independent and identically distributed across the pixels, so that $E(N_{ij}N_{kl}) = \sigma_{ij}^2\delta_{ik}\delta_{jl}$. What makes this problem more complicated than the standard linear least-squares estimate is that the noise is Poisson distributed, so the variance at pixel (i,j) is given by $\sigma_{ij}^2 = B + AP_{ij}$ (where we are ignoring the readout noise).

A1. LINEAR LEAST-SQUARES ESTIMATE

Since the background B is known, the weighted least-squares estimate is found by minimizing the following cost with respect to A:

$$J = \frac{1}{2}\sum_{ij} W_{ij}\left(I_{ij} - B - AP_{ij}\right)^2, \tag{A1}$$

where the W_{ij} are some to-be-determined weights.

This results in the optimal estimate,

$$\hat{A} = \frac{\sum_{ij}(I_{ij} - B)W_{ij}P_{ij}}{\sum_{ij} W_{ij}P_{ij}^2}. \tag{A2}$$

The optimal weights are found by forming the variance of the estimate and minimizing with respect to W_{ij} (minimum-variance estimation),

$$\sigma_A^2 = \frac{\sum_{ij} \sigma_{ij}^2 P_{ij}^2 W_{ij}^2}{\left(\sum_{ij} W_{ij}P_{ij}^2\right)^2}. \tag{A3}$$

The result is that $W_{ij} \propto 1/\sigma_{ij}^2$. That is, as is well known, the optimal weights are given by the noise variance. This implies that the cost function to be minimized is given by

$$J = \frac{1}{2}\sum_{ij} \frac{\left(I_{ij} - B - AP_{ij}\right)^2}{AP_{ij} + B}. \tag{A4}$$

This equation has theoretical appeal, since for large photon counts, the Poisson noise approaches a Gaussian distribution, and this is related to the maximum likelihood estimate. Note, however, that A appears nonlinearly in the cost, and an exact maximum likelihood solution requires a more complex calculation. We address that further below.

It is important to note that there is a bit of circular reasoning here. The expression for $\hat{A}$ in equation (A2) was found by minimizing the cost with respect to A assuming the weights were independent of A. In the standard problem formulation, this is usually the case and the resulting estimate in equation (A2) is optimal. Here, however, this result was used to find that the optimal weights were proportional to the noise variance, which for Poisson noise is a function of A! This puts into doubt the optimality of the resulting expression. However, the case of most interest here is the background-limited case in which $B \ggg AP_{ij}$. The weights then become constant and thus cancel out of the estimate. The resulting estimate is given by

$$\hat{A} = \frac{\sum_{ij}(I_{ij} - B)P_{ij}}{\sum_{ij} P_{ij}^2}, \tag{A5}$$

and the variance of the error is given by

$$\sigma_A^2 = \frac{\sum_{ij} P_{ij}^2 \sigma_{ij}^2}{\left(\sum_{ij} P_{ij}^2\right)^2} = \frac{B}{\sum_{ij} P_{ij}^2}, \tag{A6}$$

where, in this background-limited case, $\sigma_{ij}^2 \equiv B$.

A typical measure of the estimate quality is the signal-to-noise ratio given by dividing A by the standard deviation of its estimate and squaring, resulting in, for the background-limited case,

$$\left(\frac{A}{\sigma_A}\right)^2 = \frac{A^2 \sum_{ij} P_{ij}^2}{B}. \tag{A7}$$

Despite being nonoptimal for the case in which the background and planet are comparable, equation (A5) has the appeal of being unbiased; it is therefore often used as the estimate for the general case as well. It is tempting to attempt a nonlinear least-squares solution of the cost in equation (A4), hoping that although nonoptimal (the weights do not represent the minimum variance because the dependence on A was ignored), it will still provide an improved S/N. Unfortunately, this is not the case. Not only is the problem much more difficult, with no simple closed form solution, the resulting estimate is biased and the variance is larger.

A2. BLUE ESTIMATE

In his appendix, Burrows (2003) used the BLUE (best linear unbiased estimate) approach rather than least-squares minimization. In fact, this is sensible as it avoids some of the circular reasoning problems above. We reproduce that here and show that the same signal-to-noise ratio results.

The most general linear, unbiased estimate of the irradiance is given by

$$\hat{A} = \frac{\sum(I_{ij} - B)W_{ij}}{\sum W_{ij}P_{ij}}. \tag{A8}$$

It is straightforward to show that this is an unbiased estimate for any W_{ij} as long as N_{ij} is zero mean. The variance of this estimate is found via

$$E\left[\left(A-\hat{A}\right)^2\right]=\frac{\sum_{ij}\sum_{kl}\delta_{ik}\delta_{jl}\sigma_i^2W_{ij}W_{kl}}{\left(\sum_{ij}W_{ij}P_{ij}\right)^2},$$
$$\sigma_A^2=\frac{\sum_{ij}\sigma_{ij}^2W_{ij}^2}{\left(\sum_{ij}W_{ij}P_{ij}\right)^2}. \tag{A9}$$

This is the same expression as in equation (D.1) of Burrows (2003). Following the same procedure as above, the optimal weights are found by maximizing the signal-to-noise ratio or, equivalently, minimizing the variance. The resulting expression for the optimal weights is

$$W_{ij}=\frac{P_{ij}}{\sigma_{ij}^2}\left(\frac{\sum_{ij}\sigma_{ij}^2W_{ij}^2}{\sum_{ij}W_{ij}P_{ij}}\right). \tag{A10}$$

This shows that the weights are proportional to P_{ij}/σ_{ij}^2, as the term in parentheses is independent of (i,j). Since any constant term cancels out of the estimate, the weights can be set equal to P_{ij}/σ_{ij}^2 without loss of generality,

$$W_{ij}=\frac{P_{ij}}{AP_{ij}+B}. \tag{A11}$$

The circularity problem of Poisson processes shows up here when this expression for the weights is substituted into the BLUE estimate in equation (A8), as the expression for $\hat{A}$ depends on A itself. This is because the exact problem is inherently nonlinear; the linear estimate is only unbiased if the weights are forced to be independent. For an exact optimal estimate, one must turn to maximum likelihood, or other, approaches. However, again assuming a background-limited signal, where $AP_{ij}\ll B$, the weights can be taken to be $W_{ij}=P_{ij}$. Substituting this into equation (A8) results in the estimate equation,

$$\hat{A}=\frac{\sum\left(I_{ij}-B\right)P_{ij}}{\sum P_{ij}^2}. \tag{A12}$$

This is exactly the same as equation (A5) for the least-squares estimate and what we use for the PSF fitting detection in equation (20).

As before, the expression for the weights, $W_{ij}=P_{ij}$, can be substituted into the variance expression (eq. [A9] above) to form a signal-to-noise ratio equation,

$$\left(\frac{A}{\sigma_A}\right)^2=\frac{A^2\sum_{ij}P_{ij}^2}{B}. \tag{A13}$$

It is interesting to look at the opposite case in which the background is negligible. In that case, $W_{ij}\cong 1/A$, a constant that cancels out of equation (A8). This is equivalent to the least-squares solution in § A1 weighted only by the PSF. This is a very common photometric approach to PSF fitting. The estimate simply becomes

$$\hat{A}=\frac{\sum I_{ij}}{\sum P_{ij}}, \tag{A14}$$

with variance

$$\sigma_A^2=\frac{A}{\sum P_{ij}}. \tag{A15}$$

This is an unbiased estimate with S/N of $\left(A\sum P_{ij}\right)^{1/2}$.

It is important to note that the estimator in equation (A12) is still unbiased even if the background and signal are comparable; it is just no longer minimum variance. It is for that reason that it is used as the estimator for the detection problem. It is interesting, then, to compute the variance and resulting signal-to-noise ratio of the estimator in equation (A12) but with the full noise variance. The expression for the estimate variance becomes

$$\sigma_A^2=\frac{A\sum_{ij}P_{ij}^3}{\left(\sum_{ij}P_{ij}^2\right)^2}+\frac{B}{\sum_{ij}P_{ij}^2}, \tag{A16}$$

resulting in an expression for the signal-to-noise ratio including the point source photon noise,

$$\left(\frac{A}{\sigma_A}\right)^2 = \frac{A^2\left(\sum_{ij} P_{ij}^2\right)^2}{A\sum_{ij} P_{ij}^3 + B\sum_{ij} P_{ij}^2}. \tag{A17}$$

Because of the Poisson nature of the signal and the resulting dependence of the noise variance on the signal itself, the linear, unbiased estimator is no longer optimal. Nevertheless, the linearity of the equation makes performing a detection and finding integration time estimates straightforward. It is, therefore, the most common approach to PSF fitting. It is, however, appealing to investigate the possibility of a maximum likelihood estimator that accounts for this dependence. We summarize such an approach next.

A3. POISSON MAXIMUM LIKELIHOOD ESTIMATION

Since the photons are known to arrive via a Poisson process, it is sensible to turn to maximum likelihood estimation under the Poisson distribution in an attempt to solve the optimality problem.[2] For simplicity, we start by examining the zero background case. The Poisson probability distribution at pixel (i, j) is given by

$$p_{ij} = e^{-AP_{ij}} \frac{\left(AP_{ij}\right)^{I_{ij}}}{I_{ij}!}. \tag{A18}$$

The total measurement probability across all the pixels is the product of the individual probabilities, as the measurements at each pixel are independent:

$$p = e^{-\sum AP_{ij}} \frac{\prod\left(AP_{ij}\right)^{I_{ij}}}{\prod I_{ij}!}. \tag{A19}$$

Taking the logarithm gives the log-likelihood function,

$$L = -\sum AP_{ij} + \sum I_{ij} \ln A, \tag{A20}$$

where we have dropped constant terms independent of A. Taking the derivative and setting equal to zero gives the maximum likelihood estimate,

$$\hat{A} = \frac{\sum I_{ij}}{\sum P_{ij}}. \tag{A21}$$

This is identical to the BLUE estimate with zero background in equation (A14) and the PSF weighted linear least squares. For signals with little to no background, this is the most common optimal estimate of the intensity.

For the Poisson maximum likelihood with background, we form the total probability density function,

$$p = e^{-\sum AP_{ij}+NB} \frac{\prod\left(AP_{ij} + B\right)^{I_{ij}}}{\prod I_{ij}!}, \tag{A22}$$

where N is the number of pixels used in the estimate sums. We again form the log-likelihood function,

$$L = -\sum AP_{ij} + \sum I_{ij} \ln\left(AP_{ij} + B\right), \tag{A23}$$

and take the derivative and set equal to zero to find the equation that $\hat{A}$ must satisfy,

$$\sum \frac{I_{ij}P_{ij}}{\hat{A}P_{ij} + B} = \sum P_{ij}. \tag{A24}$$

[2] It is tempting to first search for a ML estimator assuming a Gaussian distribution of the noise, particularly since for large photon counts the Poisson distribution is close to Gaussian. However, Gaussian maximum likelihood in this case is difficult, highly nonlinear, and results in a biased estimate with larger error than the approximate linear cases.

For the general case of comparable background and signal, there is no closed form solution to this expression for $\hat{A}$. One must resort to numerics. However, taking the expected value of both sides does show that the estimate is unbiased. At the cost of a loss of linearity, this is the best PSF fitting estimate for A, although it is difficult to find an expression for the signal-to-noise ratio that can be used to determine integration times. Monte Carlo approaches are the most convenient.

For the background-limited case ($B \gg AP_{ij}$) it is possible to find an approximate expression for $\hat{A}$ by expanding equation (A24),

$$\hat{A} \cong \frac{B}{\sum I_{ij} P_{ij}^2} \left(\sum I_{ij} P_{ij} - B \sum P_{ij} \right). \tag{A25}$$

Again, the assumption is that the background is known. Because of the first-order expansion, this estimate has a bias of $\mathcal{O}(1/B)$, small for the background-dominated case. Because this ML expression does not lead to useful expressions of the integration time, we continue to use the (perhaps) more conservative linear approach.

REFERENCES

Brown, R. A. 2005, ApJ, 624, 1010

Brown, R. A., & Burrows, C. J. 1990, Icarus, 87, 484

Burrows, C. J. 2003, in The Design and Construction of Large Optical Telescopes, ed. P. Y. Bely (New York: Springer), 433

Kasdin, N. J., Vanderbei, R. J., Littman, M. G., & Spergel, D. N. 2005, Appl. Optics, 44, 1117

Kasdin, N. J., Vanderbei, R. J., Spergel, D. N., & Littman, M. G. 2003, ApJ, 582, 1147

Lyon, R. G., Hollis, J. M., & Dorband, J. E. 1997, ApJ, 478, 658

Vanderbei, R. J., Kasdin, N. J., & Spergel, D. N. 2004, ApJ, 615, 555

Vanderbei, R. J., Spergel, D. N., & Kasdin, N. J. 2003a, ApJ, 590, 593

———. 2003b, ApJ, 599, 686