

Point Process Algorithm: A New Bayesian Approach for TPF-I Planet Signal Extraction

T. Velusamy, K. A. Marsh and B. Ware

Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, California 91109, U.S.A.

email: velu@jpl.nasa.gov; Kenneth.A.Marsh@jpl.nasa.gov; Brent.Ware@jpl.nasa.gov

Abstract. TPF-I capability for planetary signal extraction, including both detection and spectral characterization, can be optimized by taking proper account of instrumental characteristics and astrophysical prior information. We have developed the Point Process Algorithm (PPA), a Bayesian technique for extracting planetary signals using the sine/cosine chopped outputs of a dual nulling interferometer. It is so-called because it represents the system being observed as a set of points in a suitably defined state space, thus providing a natural way of incorporating our prior knowledge of the compact nature of the targets of interest. It can also incorporate the spatial covariance of the exozodi as prior information which could help mitigate against false detections. Data at multiple wavelengths are used simultaneously, taking into account possible spectral variations of the planetary signals. Input parameters include the sigma of measurement noise and the *a priori* probability of the presence of a planet. The output can be represented as an image of the intensity distribution on the sky, optimized for the detection of point sources. Previous approaches by others to the problem of planet detection for TPF-I have relied on the potentially non-robust identification of peaks in a “dirty” image, usually a correlation map. Tests with synthetic data suggest that the PPA provides greater sensitivity to fainter sources than does the standard approach (correlation map + CLEAN), and will be a useful tool for optimizing the design of TPF-I.

Keywords. methods: data analysis, techniques: interferometric, stars: planetary systems.

1. Introduction

A key component of the proposed Terrestrial Planet Finder mission is a nulling interferometer (TPF-I) for the wavelength range 7–15 μm , currently envisaged as a free-flying four-element dual-Bracewell array. The outputs of two nulling pairs are combined coherently to produce a modulated signal as the interferometer is rotated about the line of sight to the star (Beichman & Velusamy 1999). This signal may be “sine-chopped” by differencing two phase-shifted versions ($\pi/2$ & $3\pi/2$) to maximize the information content for planetary signals (Velusamy & Beichman 2001).

Various inversion procedures exist for extracting the fluxes and locations of possible planets (see, for example, Angel & Woolf 1997; Velusamy & Marsh 2004); the one most widely used involves making a correlation map (referred to as a “dirty image”) followed by deconvolution with the CLEAN algorithm (Draper *et al.* 2005). A key step in that procedure is an iterative search for peaks in the dirty image, which can be non-robust when noise bumps fall on “side lobes” of the responses to other sources.

In this paper we propose a Bayesian technique for planet detection which avoids the noise-vulnerable peak-finding step. It can process data at many wavelengths

simultaneously, and can incorporate prior information such as the spatial covariance of the exozodi. Since the planetary signals are represented as a random set of points in a suitably-defined state space, we refer to the resulting algorithm as a ‘‘Point Process Algorithm’’ (PPA) following the terminology of Richardson & Marsh (1992).

2. Measurement Model

The starting point of our approach is a measurement model which relates a data vector, \mathbf{d} (whose components are the complete set of samples of the sine-chop signal at all wavelengths of observation), to the intensity distribution, $I(x, y)$, on the sky. The latter is modeled as the superposition of a set of point sources of unknown number, fluxes and positions, upon an extended background whose intensity at position (x_j, y_j) is denoted by ζ_j . The distribution of point sources is represented as a set of occupation numbers in a 3-dimensional state space whose axes are flux and x, y position. The state space is divided into a regularly-sampled grid of cells, such that the cell labeled (f_k, x_j, y_j) represents the k th possible flux value, f_k , at the j th spatial position, (x_j, y_j) . If a point source of that particular flux density is present at that particular position, then the occupation number of that cell, Γ_{jk} , will be equal to 1; otherwise, it will equal zero. This representation provides a convenient way of incorporating our prior knowledge of spatial dilution.

The measurement model must include a spectral model which relates the planetary fluxes at the different wavelengths. There are several possibilities, one of which is to treat the planets as black bodies at the local radiation temperature. If the orbital inclination and position angle of the orbital tilt axis are known, the planet temperature, $T_j \equiv T(x_j, y_j)$, can then be determined at each (x_j, y_j) from knowledge of the stellar luminosity and radial distance. Our measurement model can then be written as:

$$d_i = \sum_{j,k} f_k [B(\lambda_i, T_j)/B(\lambda_0, T_j)] H_i(x_j, y_j) \Gamma_{jk} + \sum_j H_i(x_j, y_j) \zeta_j(\lambda_i) + \nu_i \quad (2.1)$$

where λ_i represents the wavelength corresponding to the i th measurement, $B(\lambda, T)$ represents the Planck function, and the f_k represent possible values of planetary flux at some suitably-defined reference wavelength, λ_0 ; $H_i(x_j, y_j)$ represents the response of the measurement system to a source of unit flux at (x_j, y_j) , and ν_i is the measurement noise. We can rewrite (2.1) in more compact notation as:

$$\mathbf{d} = \mathbf{A}\mathbf{\Gamma} + \boldsymbol{\mu} \quad (2.2)$$

where $\boldsymbol{\mu}$ is a stochastic quantity representing the combined effects of ν_i and $\zeta_j(\lambda_i)$, and $\mathbf{\Gamma}$ is the state vector whose components are Γ_{jk} ; \mathbf{A} is a matrix whose elements are derived from values of $H_i(x_j, y_j)$ and other quantities as indicated by (2.1).

We assume ν_i is an uncorrelated zero-mean Gaussian random process (GRP). The exozodi distribution, $\zeta_j(\lambda_i)$ is also assumed to be a GRP, but with covariance C_ζ corresponding to the spatial correlation properties. If an estimate of the exozodi distribution is available from lower-resolution (possibly ground-based) observations, its contribution can be subtracted *a priori* from \mathbf{d} , so that $\zeta_j(\lambda_i)$ may be regarded as zero-mean. The state vector, $\mathbf{\Gamma}$, is also regarded as a stochastic process whereby each cell is statistically uncorrelated, and the *a priori* probability of cell occupancy is a small number, P_1 .

3. Estimation Procedure

The central operation is to estimate the state vector, $\mathbf{\Gamma}$, given the observed data. Specifically, we will obtain the expectation value of this quantity, defined by:

$$\rho(x_j, y_j, f_k | \mathbf{d}) \equiv E \Gamma_{jk} \quad (3.1)$$

where E is the expectation operator. We refer to $\rho(x_j, y_j, f_k | \mathbf{d})$ as a density, since it represents the average local density of occupied cells in the state space of position and flux.

Estimating the density in state space given a set of observed data is a generic problem in statistical mechanics, and has been applied to acoustical imaging (Richardson & Marsh 1987) and target tracking (Richardson & Marsh 1992). The solution procedure involves the solution of a hierarchy of integro-differential equations which fortunately can be truncated, to a good approximation, at the first member. We then obtain:

$$\frac{\partial \rho}{\partial t} + \phi_1 \rho = 0 \quad (3.2)$$

where t is a dimensionless progress variable representing the degree of conditioning on the data, and ϕ_1 is a conditioning factor which incorporates the data, \mathbf{d} , and the system response matrix, \mathbf{A} .

Equation (3.2) is to be integrated from $t = 0$ to a terminal value, t_f , defined by a stopping rule based on consistency with the observed data, corresponding to a reduced chi squared value (χ_r^2) of unity.

The initial condition on ρ for the numerical solution of (3.2) is $\rho(t = 0) = \rho_0$, where ρ_0 is the *a priori* density equal to the constant value P_1 ; its integral over all state space represents the *a priori* expectation number of planets present.

Our corresponding estimate of the source intensity distribution is then:

$$\hat{I}(x_j, y_j) = \sum_k f_k \rho(x_j, y_j, f_k) \quad (3.3)$$

Estimates of the planet fluxes themselves may be obtained from the integrated value around each peak in this image; the uncertainties correspond to the standard *a posteriori* variances of maximum likelihood estimates.

4. Tests with Synthetic Data

We have tested the PPA using a set of 15 test cases designed to replicate realistic observing scenarios for TPF-I in the wavelength range 7–15 μm . Each case involved 0–5 planets at radial distances of 0.4–5.25 AU from a solar-type star at a distance of 10 or 15 pc, superposed on a 1 Zodi dust distribution. The 15 cases involved a total of 37 planets, 24 of which were of 1 Earth flux, and 2 of which were < 1 Earth flux. The orbital inclination was 60° except for two face-on cases. All planets were at Earth temperature (260 K). Radial distances and signal-to-noise ratios were distributed as shown in Fig. 1.

The measurement configuration was the “X array” consisting of four 4-m detectors at the vertices of a rectangle of width 12 m (corresponding to the nulling baseline length) and a length of either 36 m or 72 m. Note that we have adopted 12m for nulling in lieu of the nominal design value of 20m to improve the sensitivity at the short wavelengths. The measurements consisted of the simulated sine chop signals through a full rotation of the array, at wavelengths of 7.44, 8.50, 9.92, 11.90 & 14.90 μm , and a total integration time of 1 day. The data generation included the effects of stellar leakage (due to incomplete

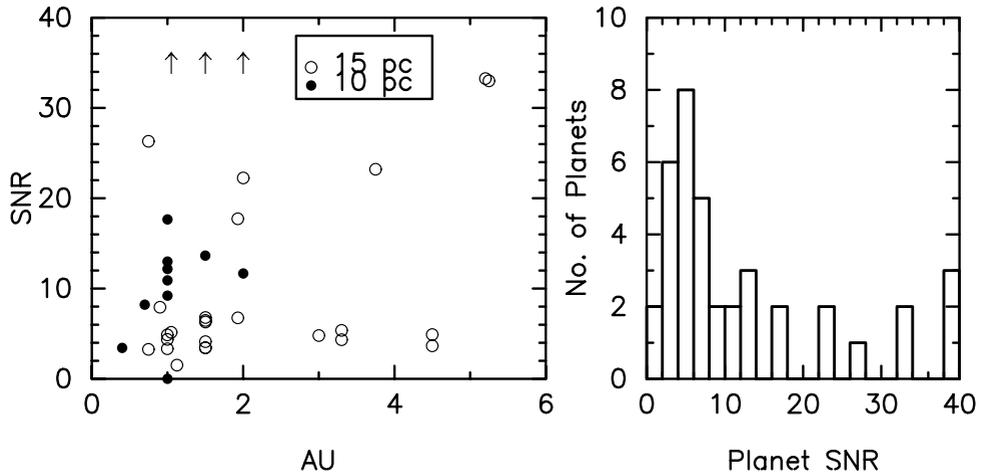


Figure 1. The distribution of planet parameters in the test cases, with respect to radial distance from the star and corresponding ideal SNR for an isolated planet. The distances to the stars are indicated.

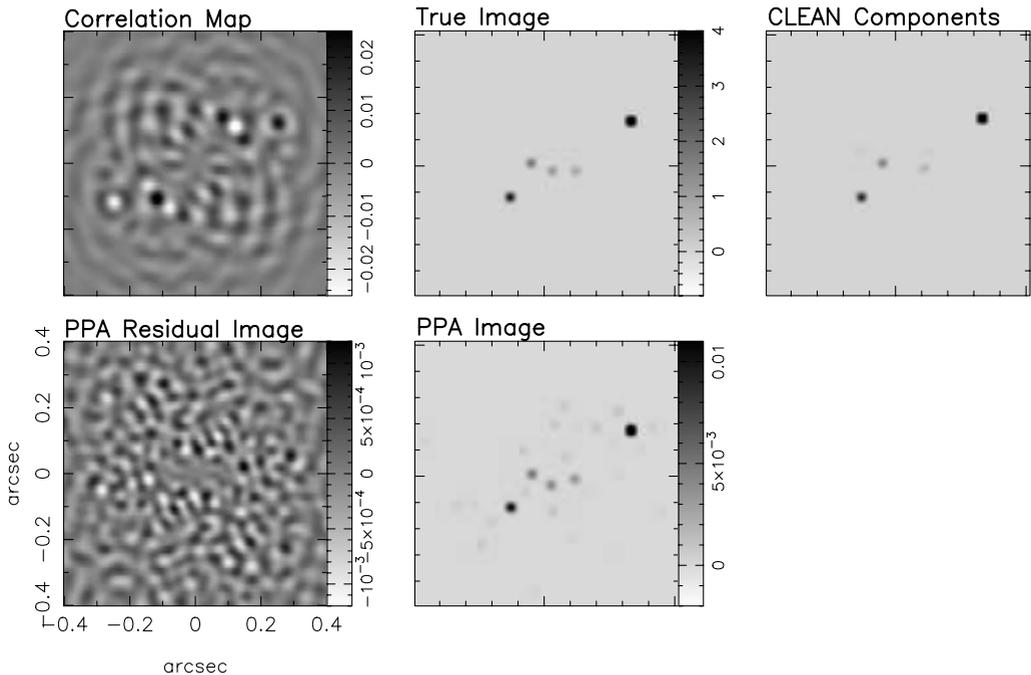


Figure 2. Images obtained by PPA and Clean methods for a test case involving 5 planets with the X-36m array. Similar images were obtained for all 15 test cases.

nulling of the stellar photosphere) and thermal emission from a symmetrical exozodiacal dust cloud. The noise model included the effects of Poisson noise from all components.

The data were inverted using the PPA, and also by the standard technique of “correlation map + CLEAN” for comparison purposes. The spatial covariance of the exozodi was ignored (i.e., C_{ζ} was set at zero) since symmetrical sources cancel in the sine chop signal. Full account was, however, taken of the Poisson noise contribution of this component. Figure 2 shows the results for a case involving 5 assumed planets.

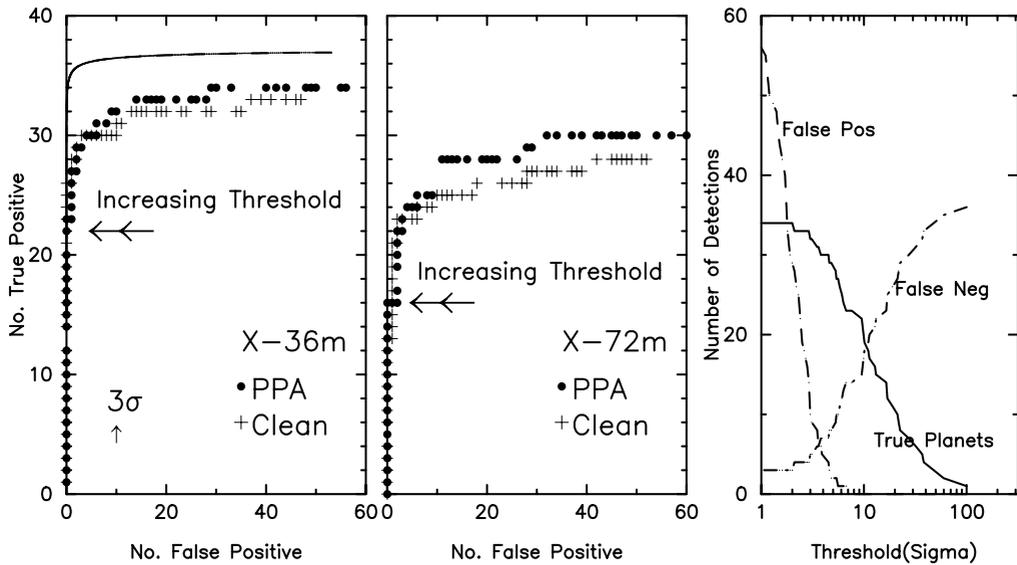


Figure 3. Planet detection statistics. The left two panels show the operating characteristic curves for PPA and CLEAN. The solid line on the left hand plot represents the idealized theoretical curve for an isolated planet. The right panel shows the detection performance of PPA as a function of threshold in sigmas.

These inversions yielded images for all 15 test cases (as exemplified by Figure 2) from which we extracted estimates of point source locations, fluxes, and corresponding uncertainties. We identified the six brightest peaks in each image, then applied a SNR threshold and counted the detections. Comparison with the true (assumed) planet positions then led to the number of true detections and false alarms as a function of the detection threshold in sigmas. Figure 3 shows a plot of true positives vs. false negatives, known as the operating characteristic curve of the detection system, for both PPA and CLEAN. It is apparent that for a large range of abscissa values, the PPA detected more planets than CLEAN, particularly with the 72 m array. The PPA performance as a function of detection threshold in sigmas is shown in Figure 3.

5. Discussion

The PPA is theoretically a near-optimal approach, and the present results support this. For the 72 m array in particular, the PPA has effectively increased the sensitivity of the array for planet detection over standard techniques exemplified by CLEAN.

The results also indicate that the 72 m array would be less sensitive than the 36 m array for this ensemble of planets, presumably due to the smaller relative range of spatial frequencies involved. By combining the data from both configurations (with same total integration time), we found that the added spatial frequency coverage provided by the 72 m array did not offset the reduced signal to noise of the 36 m data.

The detection statistics from Figure 3 indicates that:

- At the 5.8σ level we detect 26 of the 37 planets, but obtain 11 false positives; the 26 detections include 15 of the 24 Earth flux planets.
- A threshold of 5σ would facilitate detection of more than half of the Earths in a TPF survey, with only one false positive.
- Lowering the threshold to 4σ would facilitate detection of 19 of the 24 Earths in the sample, but with a penalty of 5 false positives.

Future work will emphasize the effects of exozodi, including the effects of spatial inhomogeneities which may masquerade as planets in the interferometer signals, and their mitigation via appropriate incorporation of *a priori* estimates and spatial covariance. For the latter purpose we will make use of recent observations of debris disks using the Spitzer Space Telescope (Beichman *et al.* 2005).

Acknowledgements

This work was performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

- Angel, J. R. P. & Woolf, N. J. 1997, *ApJ* 475, 373
 Beichman, C. A. *et al.* 2005, *ApJ* 622, 1160
 Beichman, C. A. & Velusamy, T. 1999 in: S. C. Unwin & R. V. Stachnick (eds.), *Optical & IR Interferometry from Ground and Space* (ASP Conference Series) Vol. 194, pp 405–419
 Draper, D. W. *et al.* 2005, *AJ*, submitted
 Richardson, J. M. & Marsh, K. A. 1987, *Acoustical Imaging* 15, 615
 Richardson, J. M. & Marsh, K. A. 1992, in: C. R. Smith *et al.* (eds.), *Maximum Entropy and Bayesian Methods* (Kluwer Academic Publishers, Netherlands), pp 213–220
 Velusamy, T. & Beichman, C. A. 2001, *IEEE Aerospace Conference Proceedings* Vol 4, p 2013
 Velusamy, T. & Marsh, K. A. 2004, *Proceedings of 2nd TPF/Darwin International Conference (Dust Disks and the Formation, Detection and Evolution of Habitable Planets)* San Diego, July 28–29, 2004

Discussion

KALTENEGGER: Do you find that the use of two baselines improves the results such that it justifies the double integration time needed then?

VELUSAMY: We find that for a given total integration time, it is better to put all of it into the 36 m configuration rather than splitting it between 36 m and 72 m, i.e. the added (u, v) coverage provided by the 72 m data did not quite compensate for the reduction in signal to noise of the 36 m data.

KALTENEGGER: Do your calculations take the properties of the instrument over a full rotation into account, i.e., do you extract the signal over a full rotation of the instrument, or over a fraction of a full rotation, or do you assume no rotation for your calculation?

VELUSAMY: The inversion makes simultaneous use of the data at all rotation angles of the array, from 0° to 360° in increments of typically 1°, as in our simulations. However the algorithm can work on any given data set with partial or complete rotation.