Science Operations and the Minimum Scale of TPF

Robert A. Brown
Space Telescope Science Institute

Stuart B. Shaklan, and Sarah L. Hunyadi
Jet Propulsion Laboratory, California Institute of Technology

1. INTRODUCTION

For any mission, the most stringent requirement is the one that determines its minimum scale—its minimum size, complexity, and cost. In recent years, it was widely assumed that *search completeness*[1] was *TPF*'s most stringent requirement, and therefore its greatest technical challenge was ensuring that the instrument design reliably detects and resolves extrasolar planets in the presence of unsuppressed starlight and zodiacal emission. These assumptions were held for both the thermal-infrared (interferometric; *TPF-I*) version and reflected-light (coronagraphic; *TPF-C*) version of the *TPF* mission. Recently, however, mission studies performed for the *TPF-C* Design Reference Mission (DRM)[2] have cast doubt on these assumptions. It now appears that the most stringent requirements of both *TPF-C* and *TPF-I* are to acquire critical astrometric information about the motion of a candidate planet in the epoch of its discovery. These new requirements are *science operational*: (R1) to ensure that a candidate planet is promptly discriminated from possible background confusion[3] and (R2) to ensure a planet's recovery at later observing epochs.[4] Until the implications of requirements R1 and R2 are understood, the minimum scale of both versions of *TPF* will remain unknown.[5] The analysis is now underway.

The most critical aspects of technical performance for meeting the new requirements are astrometric accuracy and the pointing restrictions—notably the solar avoidance angle. A timed series of measurements of a source's apparent position relative to the star must be obtained, and the data quality—accuracy and orbital coverage—must be sufficiently high to discriminate background sources and provide orbital solutions useful for predicting the future observability of the planet. For *TPF*, obtaining high-quality data is especially challenging, because (1) much of the typical planetary orbit is not available to observation,[6] (2) the intervals of time over which any star can be viewed are limited to a fraction of a year by solar avoidance, and (3) aliasing occurs for planetary periods approximately resonant with one Earth year. Because of these challenges, the astrometric accuracy of each positional measurement must be even higher than otherwise—but *how* high has not yet been determined. Because higher astrometric accuracy can only be achieved by reducing the size of the point-spread function—i.e., increasing the size of the instrument, which is the design parameter with the steepest cost slope for an instrument in space—it is important to resolve this issue decisively and correctly as soon as possible.

If the astrometric accuracy of *TPF* were insufficient to satisfy R1 and R2, then errors and inefficiencies would be introduced into science operations. Mission time would be wasted on void or misguided observations, and the

---

[1] Brown, R. A. 2005, "Single-Visit Photometric and Obscurational Completeness," *ApJ* 624: 1010–1024. Hunyadi, S. L., Shaklan, S. B., and Brown, R. A. 2007 "Single Visit Completeness Optimization," submitted to *ApJ*.

[2] Brown, R. A., Hunyadi, S. L., and Shaklan, S. B., 2006, "A DRM for *TPF-C*: A Design Reference Mission for the Coronagraphic T*errestrial Planet Finder*." http://sco.stsci.edu/tpf_downloads/TPF-C_DRM.pdf/.

[3] Brown, R. A., 2007, "Differentiating Extrasolar Planets from Background Confusion by Apparent Motion," submitted to *ApJ*, http://sco.stsci.edu/tpf_downloads/background_confusion.pdf/.

[4] DRM §5.2, p. 49.

[5] A related, existing requirement is in the *TPF-C* Science Requirements Document (SRD): "For at least 50 percent of detected planets whose projected semimajor axes lie within the habitable zone of their parent star, *TPF-C* shall measure the actual semimajor axes to within 10%. *TPF-C* shall also measure their orbital eccentricities to an absolute accuracy of ±0.3." This requirement is scientific and not time critical. It is less demanding than #2 and would be subsumed by it.

[6] DRM §5.3, p. 51.

science goals would be at risk, and possibly unobtainable. While the impacts are not yet fully and quantitatively known, they seem likely to strongly affect the design process for *TPF-I* and *TPF-C*.

In the remainder of this paper, we illustrate the science-operational issues for *TPF* with particular reference to the FB-1 design and associated DRM research for *TPF-C*. For *TPF-I*, the results should be qualitatively similar to those for *TPF-C*, involving a different planetary phase function, but similar instrumental parameters.

We show illustrative results for three typical target stars; the same results for all 136 top *TPF-C* target stars are available at *StarVault*.[7]

Two relevant discussions are available online: (1) the roles of technical performance in selecting target stars[8] and (2) the considerations for determining the photometric orbit in reflected starlight.[9]

2. THE SCIENCE-OPERATIONAL VIEWPOINT & MISSION MODELING

In the future science operations center for *TPF*, decisions about what star to observe next, and for what purpose—searching, background discrimination, orbit determination, spectroscopy, etc.—will be made continually. *Making and implementing sound scheduling decisions is the essence of science operations.* For the possible newly discovered planets, these decisions will be made with limited information available—but many possibilities about the nature of an object and how it might move and change in brightness. Because each imaging observation typically demands days, or even weeks, of precious observing time (spectroscopic observations may take ten times as long), scheduling decisions must be of the highest possible quality, making the best use of available knowledge (both observational and theoretical), while seeking more information.

Consider the aftermath of the very first *TPF* search observation of the very first target star. Assume it finds a source with a brightness and separation consistent with being a planet of interest. What's next? An immediate follow-up observation with a roll of the telescope aperture around the line of sight will confirm the external reality of the source—i.e., that it is not a speckle. But it could still be a background confusion source: very little is known about background sources as faint as those sought by *TPF*. To discriminate a possible planet from background, another observation must be made after a delay that is sufficiently long for the planetary movement (probabilistic) to be measurably different from the apparent motion of a background source induced by the parallactic and proper motion of the star (deterministic). Yet, this delay must also be sufficiently short to (1) minimize the probability that the planet has become undetectable, and (2) ensure the star is still viewable according to solar-avoidance constraints (see reference in footnote #3 for a full discussion.) Assuming the discriminating observation has confirmed the companionship of the source, additional astrometric observations must be made—in the discovery epoch, before the annual viewability window closes—to determine the orbit well enough to ensure future recovery. If we don't know the orbit at the time the viewability window closes, then at least six months will pass as the solar-avoidance restriction lapses, and we can only guess about when the companion might be observable again. The odds of a successful guess will be different for each star, but rarely high, as we will see §3.4. And if the planet is *not* recovered, we can only guess again and again, potentially wasting much more time, or give up, declaring the planet lost.

If this logic for the very first observation seems complex and risk riddled, consider how much more daunting the situation will be later in the mission, with many stars and possible planets in play at various levels of knowledge, including the possibility of multiple planetary systems, and with many possible observations all competing to be the next observations. The stakes are high, because mission time is limited and exquisitely valuable. Currently, NASA plans for this mission to last only five years.

---

[7] http://sco.stsci.edu/starvault/.

[8] http://sco.stsci.edu/tpf_downloads/on_target_stars.pdf/.

[9] Brown, R., 2006, "Determination of the Photometric Orbit in Reflected Starlight," http://sco.stsci.edu/tpf_downloads/photometric_orbit.pdf/.

Fortunately, it is possible to achieve the science-operational viewpoint and test a mission design by Monte Carlo simulation and mission modeling. Using double blind experimental formats, we can test the ability of a design to perform the necessary observations in the restricted informational environment of the mission.[10] In particular, we can use Monte Carlo simulations to estimate mission success as a function of astrometric accuracy. Mission modeling is also useful for testing the decision-making process itself. For example, modeling can quantify (in terms of probabilities) the consequences of *not* discriminating background, or of simply guessing about future observability. The ultimate goal of mission modeling is confirming that an optimized *combination* of two designs—one for the instrument and one for science operations—meets the science requirements for the mission.

3. EXEMPLIFICATION OF SCIENCE-OPERATIONAL ISSUES FOR *TPF-C*

The *TPF-C* DRM defines particular models of stellar qualification,[11] planets of interest (habitable zone),[12] and instrumental performance,[13] which have been used for the calculations reported here.

In assessing the suitability of target stars, an important quantity is the discovery rate R, for which convenient units are milliplanets per hour. For the initial search observation, $R_1 = C_1/T_{EXP}$, where C_1 is the search completeness (the fraction of Earth-like planets in habitable-zone orbits that are detectable at the time of the first observation), and T_{EXP} is the exposure time to detect a source at the limiting delta magnitude (Δmag_0). C_1 is always less than one, because planets can hide behind the obscuration or show the dark hemisphere to the observer. (C_∞ is ultimate completeness, achieved by exhaustive searching, using many searching observations at staggered times. C_∞ may equal one, but also may be less, as some possible planets may be permanently hidden behind the coronagraphic mask.)

Table 1. Three Typical *TPF-C* Target Stars

Star	HIP 23311	HIP 79672	HIP 61174
Ranking	71st of 136	68th of 136	59th of 136
Spectral Type	K4III	G5V	F2V
Apparent magnitude m_v	6.22	5.49	4.29
Absolute magnitude M_v	6.50	4.75	2.99
Color $B–V$	1.05	0.66	0.38
Parallax Π (mas)	113.5	71.3	54.9
Luminosity (relative to Sun)	0.14	0.95	5.42
Mass (relative to Sun)	0.86	1.10	1.46
Δmag_0 (optimized)	23.5	25.0	25.0
Exposure time T_{EXP} (hrs)	28	114	17
C_1	7%	35%	6%
R_1 (milliplanets hr^{-1})	2.59	3.11	3.34
Total discovery fraction C_∞	36%	94%	39%
$T_{10\%\ LOST}$ (days)	2.5	9.5	7.9
F_{DISC}	49%	45%	40%
F_{SOLAVD}	47%	47%	47%
$F_{SOLAVD \cap DISC}$	12%	1%	40%

[10] For one example of Monte Carlo mission simulation and testing, see: Brown, R. A. 2006a, "Expectations for the Early *TPF-C* Mission," in *Direct Imaging of Exoplanets: Science & Techniques*, Proceedings IAU Colloquium No. 200, (C. Aime & F. Vakili eds.), http://sco.stsci.edu/tpf_downloads/mission_studies_appendices/MS-1_Brown_Mission.pdf/.

[11] DRM §2.1, p. 5.

[12] DRM §2.2, p. 11.

[13] DRM §2.3, p. 13.

We focus on three typical stars with similar rankings (similar discovery rates R_1), but which offer quite different science opportunities and science-operational challenges. Various relevant stellar parameters are shown in Table 1. We use these three stars to exemplify four features relevant to science operations for *TPF*: (1) completeness and discovery rate; (2) properties of the discoverable subset of planets; (3) discrimination of background confusion by testing for common proper motion; and (4) planetary recovery at future observing epochs.

3.1 Initial Search Completeness, C_1, and Discovery Rate, R_1

We assume the maximum detectable magnitude difference for *TPF-C* is $\Delta mag_0 = 25$ (a factor of 10 billion), set by uncontrolled speckles of unsuppressed starlight, which can masquerade as planets. The angular radius of the central obscuration is assumed to be 57 milliarcseconds (mas), which means the minimum apparent separation of an observable planet is $s_0 = 57/\Pi$au, where Π is the annual parallax of a star in mas.

We introduce our Monte Carlo techniques for mission modeling by discussing the estimate of C_1. As illustrated in Figure. 1, we prepared random samples of 10,000 planets for each star, drawing them one at a time from the planetary population of interest. (In the computer, we use random deviates that emulate the specified probability distributions of orbital and planetary parameters to produce particular values of the randomized orbit for each planet.) Then, we counted the number of these random planets that can be detected—i.e., tested and found those that satisfy the constraints $s > s_0$ and $\Delta mag < \Delta mag_0$, where s is the apparent separation between star and random planet, and Δmag is the ratio of planetary to stellar flux expressed in magnitudes. Our estimate of C_1 was then the number detectable, divided by 10,000. (We estimated C_∞ by repeating this "observation" in the computer many times at random epochs—allowing the unfound planets to revolve in their orbits and evolve in their brightness—and counting planets as they were found.)

In mission *simulations*, the number of planets actually found in an initial search, $n_p = 0$ or 1, is a Bernoulli random variable with probability C_1.

Table 1 shows that the rank ordering of the three typical target stars tracks the value of discovery rate R_1, as expected.[14] The values of C_1 differ by a factor 5–6, which is compensated by a similar variation in (optimized) exposure time T_{EXP} (0.7 to 4.7 days).

3.2 Discoverable Subset of Planets

Each dot in Figure 1 is a particular planet of interest, possessing unique—and known—orbital elements: period, semimajor axis, three orientational (Euler) angles, and orbital phase (mean anomaly) at the initial observing epoch. Each *red* dot is one such planet that is also *detectable*, and therefore a member and exemplar of the discoverable subset. Knowing the orbital elements for each Monte Carlo planet, we can compute the true planetary position, the apparent planetary position (apparent separation and position angle), and apparent brightness of the planet *at any epoch*.

We can explore the behavior of the discoverable subset in various ways, for various purposes. We can study them like *real* planets, even adding noise to measurements of their position and brightness. We can use them to simulate actual observations to test our science-operational concepts, or to calculate properties of the ensemble. In particular, we can use them to test whether or not requirements R1 and R2 can be satisfied by a mission design. (See the reference in footnote #3 for a detailed approach to R1.)

[14] See *StarVault* or Table 6, DRM p. 19, for results on all 136 top *TPF-C* target stars.

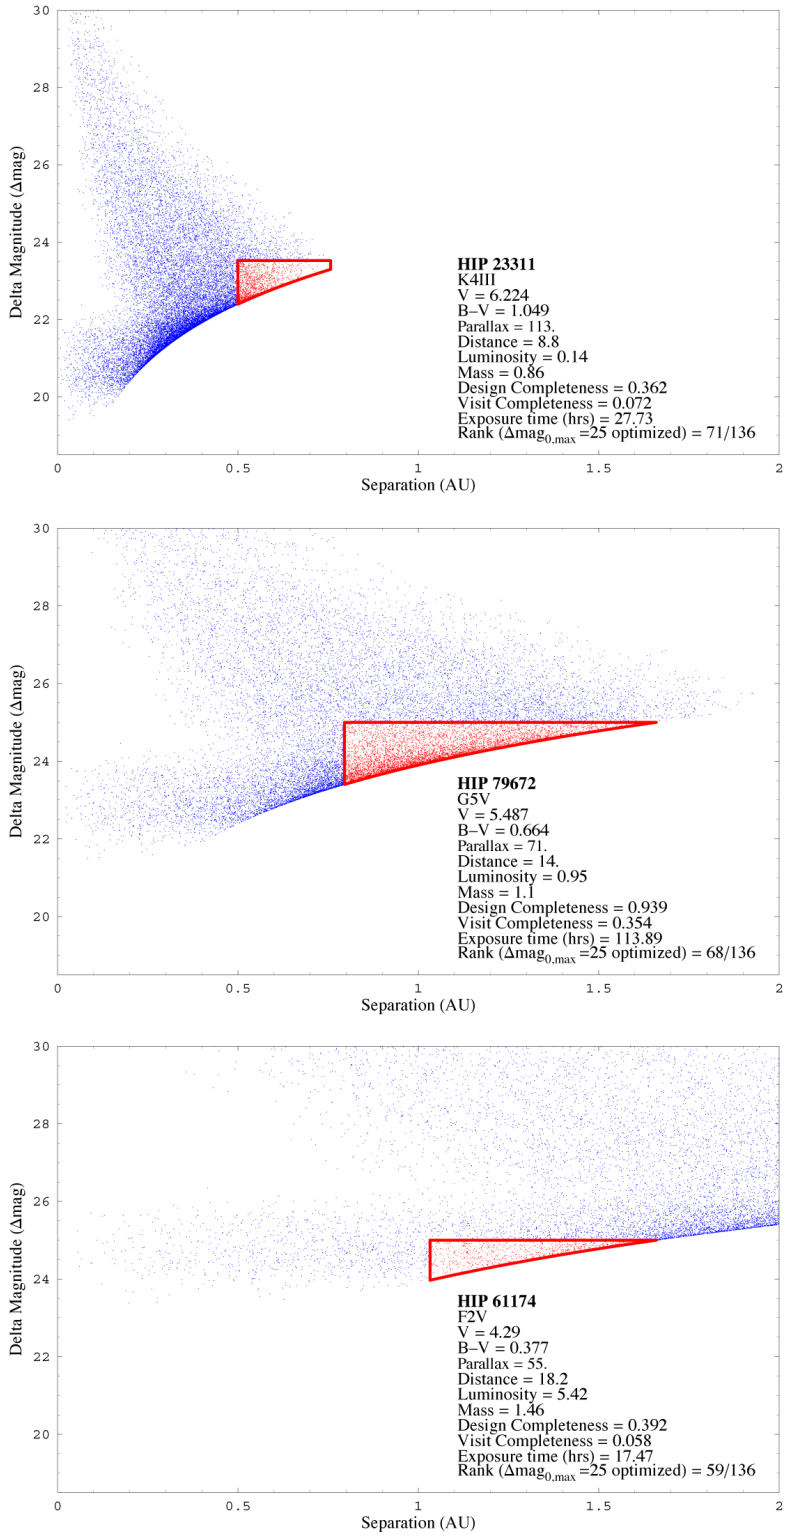

Figure 1: *Estimating C_1, the fraction of planets discovered. Out of 10,000 random planets from the population of interest, the red dots are members of the "discoverable subset" and the blue dots are not detectable. Instrumental (s_0, Δmag_0) and stellar (L, Π) factors select different discoverable subsets for each star.*

In simulations of observations, for example, we might proceed as follows. Let η_{Earth} be the probability that a star has a planet of interest. The number of planets and confusion sources actually found in the next search, $n_p = 0$ or 1, is a Bernoulli random variable with probabilities η_{Earth} C. In each trial, we would generate a sample value of n_p from a Bernoulli random deviate. If $n_p = 0$ in a trial, the experiment would proceed to consider another star (after a delay to allow the remaining possible planets to evolve in brightness and position, this star could be searched again). If $n_p = 1$ in a trial, we would select one "red-dot" planet at random to be the sample planet we had found. Only the computer would know the orbital elements, of course, but it can produce simulated observations of this planet as requested for double blind simulations of the science-operational decision-making process. Also, a new sample of planets consistent with an observation could be generated—the "found subset"—as a refinement over the discoverable subset.

The general behavior of the discoverable subset (or alternatively the found subset) can be studied as follows. For example, we want to know their "depletion time." That is, if a planet of interest were found around a given star, we want to know how long could a follow-up observation be delayed before the risk is greater than, say, 10% that the planet will have moved or changed in brightness enough to become undetectable? Call this time $T_{10\%LOST}$. To address this question, we prepared 10,000 discoverable planets for each of the 136 top TPF-C target stars, and tracked them until they became undetectable, either by disappearing behind the central obscuration or by becoming fainter than Δmag_0. The results are shown in Figure 2, and the values of $T_{10\%LOST}$ for our example stars are given in Table 1 (and in *StarVault* for all 136 stars).

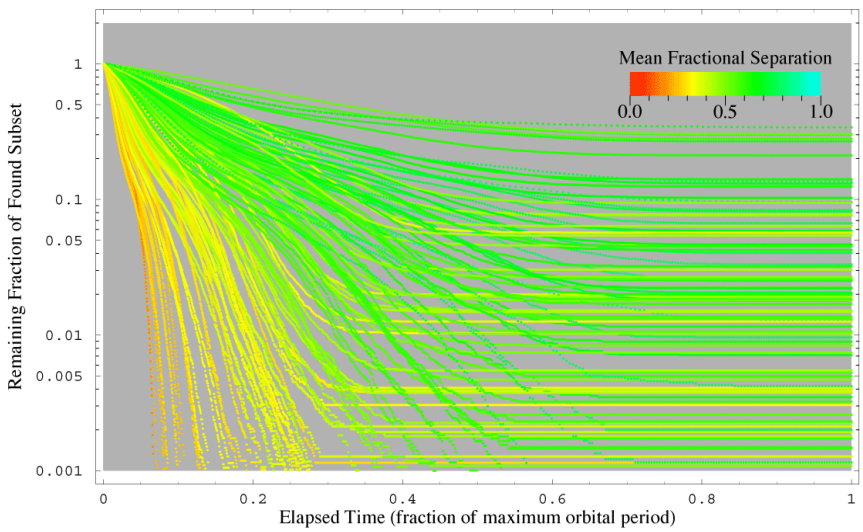

Figure 2. *The depletion with time of the discoverable subset of planets. Each curve is one star on the target list, color coded by the mean separation of the discoverable subset at the time of discovery expressed as a fraction of the maximum separation for that star (semimajor axis times one plus the eccentricity). The asymptotic value is the permanently observable fraction of the discoverable subset. The color trend is due to greater apparent transverse speed when smaller separations are preferentially sampled. The depletion curves for individual stars can be found at StarVault.* $T_{10\%LOST}$ *is the time for a depletion curve to drop to 0.90. Values of* $T_{10\%LOST}$ *are given for the three sample stars in Table 1, and for all 136 target stars in StarVault.*

In another example, we want to compute the apparent displacements of planets in the discoverable subset after some delay time, say $T_{10\%LOST}$, from the epoch of their discovery. We want this information to assess the ability of given astrometric accuracy to discriminate background confusion. This involves comparing the probability distribution of planetary displacements with the apparent displacement of a background source induced by proper motion and annual parallax.

3.3 Discriminating Background Confusion Sources[15]

It is critical to quickly confirm that a source is a planet and not background confusion. The timing of the confirming observation is a compromise between waiting too long, which risks losing the planet by "depletion," and waiting too little, which risks not achieving planetary displacements that are usefully different from the induced apparent displacement of a background source. A good compromise is $T_{10\%\mathrm{LOST}}$. Because we know the orbital elements for each planet in the discoverable (or found) subset, we can compute the probability distribution of the displacements of discoverable (or found) planets in time $T_{10\%\mathrm{LOST}}$.

By comparing the background and planetary displacements to the astrometric error of the instrument, we can estimate the probability, $1–P_{\mathrm{CONF}}$, that an observation will allow us to reject the null hypothesis, that the putative planet is a background source. For the common proper motion test to be useful, P_{CONF} must be small—less than 0.0027, say, for three standard deviations of confidence.

The solar-avoidance restriction means that stars cannot be viewed if they lie within some specified angle of the Sun (95° for *TPF-C* FB-1). This angle dictates F_{SOLAVD}, the fraction of the year that a given star can be viewed. ($F_{\mathrm{SOLAVD}} \leq 0.47$ for *TPF-C* FB-1.)

The intersection of the two conditions—stellar viewability and background discrimination utility—is of special interest. Searching observations must occur during the annual fraction when both are possible, which will be true for a fraction of the year we call $F_{\mathrm{SOLAVD}\cap\mathrm{DISC}}$. Note in Table 1 the huge difference in this quantity for the three target stars being used as examples. Figure 3 shows the details of how viewability and background-discrimination utility vary through the year for the example stars.

StarVault gives the same numerical and graphical results for all 136 *TPF-C* target stars assuming *TPF-C* FB-1. Background discrimination is *never* useful for 10% of the 136 stars, due to inadequate parallactic and proper-motion displacements, independent of solar avoidance ($F_{\mathrm{DISC}} = 0$). It is not useful for 33% of the 136 stars at the same time the star is viewable according to solar avoidance ($F_{\mathrm{SOLAVD}\cap\mathrm{DISC}} = 0$). Furthermore, $F_{\mathrm{SOLAVD}\cap\mathrm{DISC}}$ >20%, >30%, and >40% for only 50%, 41%, and 20% of the stars, respectively. Because these 136 stars are the most productive qualified stars for *TPF-C*, replacement stars will be less productive than the least productive star currently on the list. The 137th star is HIP 88175 (zeta Ser), which offers the measly productivity of 0.72 milliplanets per hour, assuming it has one planet of interest. Smaller astrometric error (larger aperture) would improve the situation.

3.4 Planetary Recovery after the Viewability Epoch of Discovery

To expose issues of planetary recoverability, we prepared 10,000 planets in the discoverable subset for each of the 136 *TPF-C* target stars. We assigned each planet a day of discovery in the calendar year that was uniformly distributed over the first 25% of the stellar viewability window for that star, with the rationale that search observations should occur as early as possible, to allow time for follow-up observations before the viewability window closes. Next, for each planet, we ran time forward, after the closing of the viewability window of the discovery epoch, until the planet became observable again (i.e., the star was viewable and the planet was detectable); this time delay, measured from the time of discovery, is the *x*-coordinate of the plots in Figure 4. Next, the time was run forward again, to the point where the planet became unobservable again, either because the star was no longer viewable or the planet was no longer detectable; this duration, measured from the epoch of the first possible recovery, is the *y*-coordinate of the plots in Figure 4.

A planet's recovery window begins at the *x*-coordinate of a dot and extends to the right for a length in days equal to the y-coordinate of the dot, as illustrated by the red lines in the middle plane of Figure 4. The extended caption to Figure 4 is a tutorial. The following is a limited commentary on the results for the three example stars.

[15] See the reference in footnote #3 for a detailed discussion of this topic.

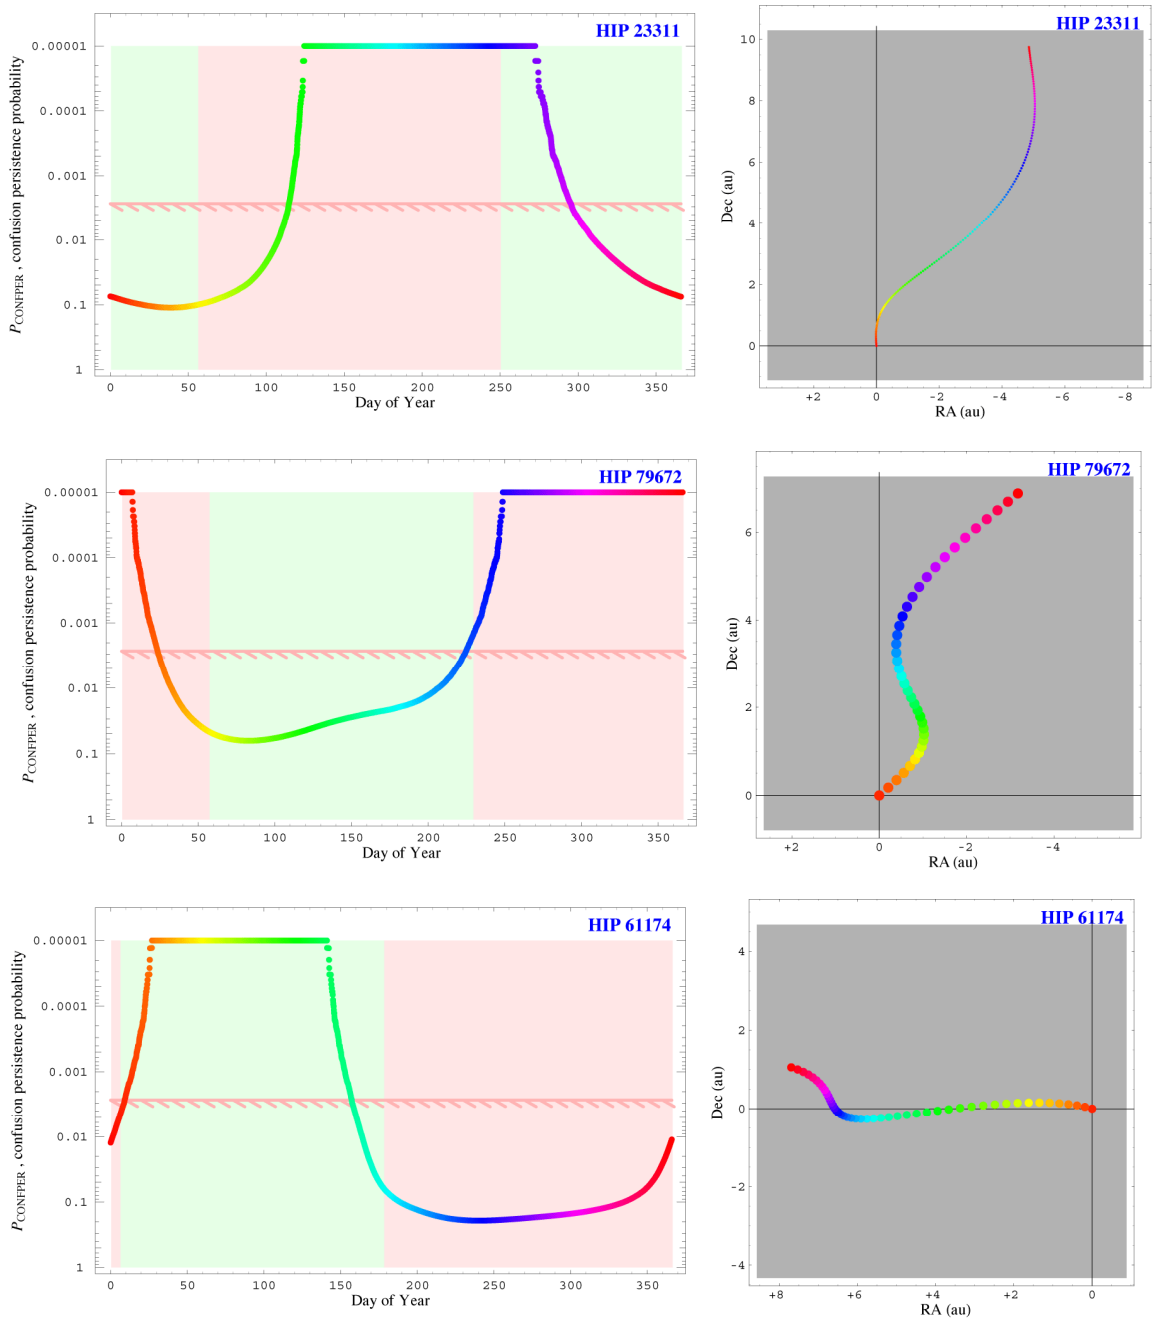

Figure 3. *Stellar viewability and background confusion differentiation. Left, light green: the days of the year when the star is more than 95° from the Sun and viewable by TPF-C FB-1. Rose: not viewable. Testing for common proper motion is useful when the curve is above the red, hatched line, which signifies three standard deviations of confidence that confusion will not persist after the second, differentiating observation. Right: displacement of background sources due to parallax and proper motion through the solar year. Points occur at intervals of* $T_{10\%LOST}$*, and the point radius is the mean amplitude of the planetary displacements in that time interval. The more widely separated the points on the right, the more useful is the confusion test. Rainbow color codes are synchronized: the larger the gap between points (right), the lower the probability of confusion persisting,* $P_{CONFPER}$ *(left). See the reference in footnote #3 for the details behind these calculations.*

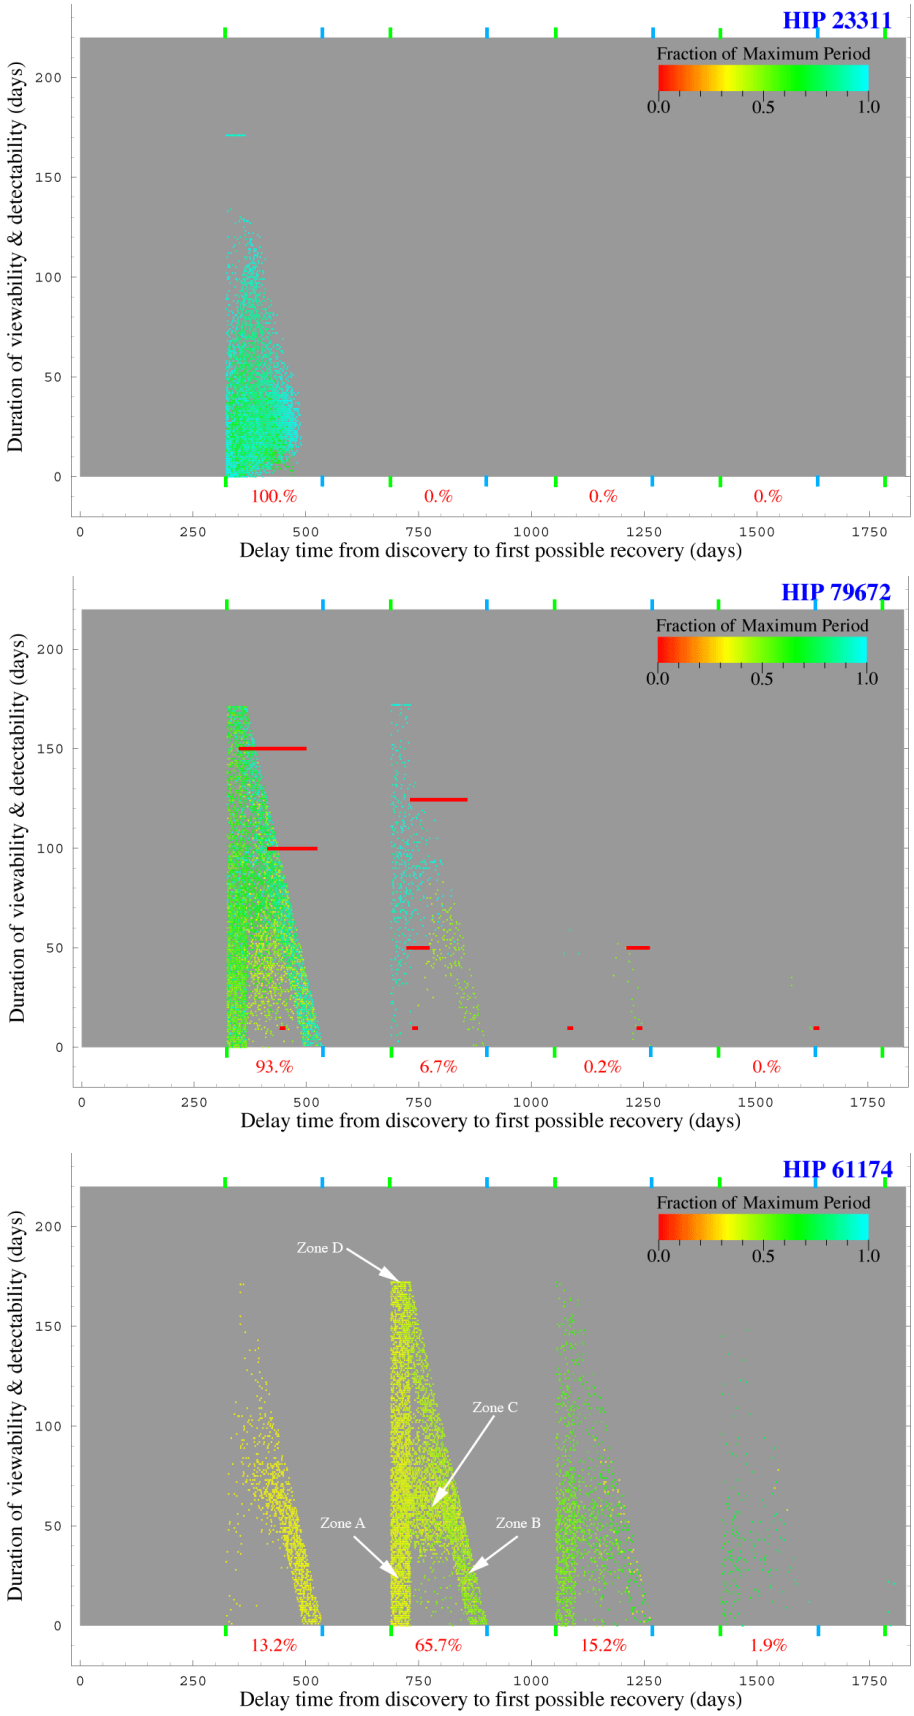

Figure 4. *(Extended figure caption follows.)*

HIP 23311 is the low-luminosity K star. From Figure 2, we see the discoverable subset consist largely of wider orbits (longer periods, as confirmed by the color code in Fig. 4), with the planets discovered near elongation. The maximum orbital period for planets of interest is 166 days, which is less than the stellar viewability window. Therefore, the entire discoverable subset is recoverable in the first stellar viewability window after the epoch of discovery. The duration of observability tends to be short, however, except for the permanently detectable planets, in orbits viewed nearly face on (the bar of dots at $y = F_{\mathrm{SOLAVD}}$ 365.25 = 172 days).

HIP 79672 is a G star, like the Sun. From Figure 2, we see the discoverable subset consist of a wider range of orbits, with the planets viewed from half to gibbous phase. The maximum orbital period for planets of interest is 616 days, which is much longer than the stellar viewability window. Some 93% of the discoverable subset is recoverable in the first stellar viewability window, and the remainder are aliased out to subsequent stellar viewability windows. The duration of observability covers the whole range, from one day to the full width of the viewability window. There are some permanently detectable planets.

HIP 61174 is a high luminosity F star. From Figure 2, we see the discoverable subset consist of a wide range of orbits viewed mostly side-on, with the planets discovered rushing to or from full phase at superior conjunction. The maximum orbital period for planets of interest is 1973 days, which is much longer than the stellar viewability window. Only 13% of the discoverable subset is recoverable in the first stellar viewability window, and these planet tend to appear late in the stellar viewability window. Some 66% and 15% can only be recovered in the second and third windows, respectively.

Planetary recoverability. *X*-axis: delay time from planetary *discovery* to the first possible *recovery* in a subsequent observability period (star viewable and planet detectable). *Y*-axis: duration of the observability period of recovery. Each dot is one of 10,000 planets representing the discoverable subset. Dot color gives the planetary period as a fraction of the longest period (365.25 days times $(a_{\max} = 1.5\ \sqrt{luminosity})^{1.5}\ (mass)^{-0.5}$, where stellar luminosity and mass are given in Table 1. The red lines on plot for HIP 79672 show the observability periods of ten particular planets explicitly. For this experiment, the time of planet discovery was confined to the first quarter of the stellar viewability period. Numbers in red give the percentage of planets that are first recoverable in each stellar viewability period.

Each stellar visibility window comprises four zones, identified in the plot for HIP 61174.

Zone A is the vertical strip that starts at $x = 365.25\ (n – 0.25 \times F_{\mathrm{SOLAVD}})$ days, for $n = 1, 2, 3...$ (green marks). This value of x is the minimum possible delay in the nth viewability window. It implies a planet that was discovered at the last possible moment in the first quarter of the initial viewability window and is first observable again at the very first possible moment of the nth subsequent viewability window. Hence the minimum possible delay of 0.75 x 365.25 days. The end of the Zone A strip is 0.25 x F_{SOLAVD} x 365.25 days later, at 365.25 x n days. It implies a planet that was discovered at the very first possible moment in the initial viewabiliity window and is observable again at the very first possible moment of the nth viewability window.

In Zone A, the observability period for the recovery of any planet is punctuated at the start by the stellar viewability window opening, and at the end by the planet becoming undetectable—by the planet going behind the central obscuration or becoming fainter than the sensitivity limit—not by the end of the stellar viewability window.

Zone B is the diagonal strip that starts at the top of Zone A and extends down to the right. At $y = 0$, Zone B ends at $x = 365.25\ (n + F_{\mathrm{SOLAVD}})$ days (blue marks). This value of x is the maximum possible delay in the nth viewability window. It implies a planet that was discovered at the first possible moment in the first quarter of the initial viewability window and is first observable again at the very last possible moment of the nth subsequent viewability window. Hence the maximum possible delay in the first subsequent viewability window of $(1 + F_{\mathrm{SOLAVD}})$ x 365.25 days. The start of the Zone B strip at $y = 0$ is 0.25 x F_{SOLAVD} x 365.25 days earlier, at $x = 365.25\ (n + 0.75 \times F_{\mathrm{SOLAVD}})$ days. It implies a planet that was discovered at the very last possible moment in the first quarter of the initial viewability window and is observable again at the very last possible moment in the nth observability window.

In Zone B, the observability period for the recovery of any planet is punctuated at the start by the planet becoming detectable, and at the end by the stellar viewability window closing.

Zone C is the region under Zones A and B. Here, the observability period for the recovery of any planet is punctuated at the start by the planet becoming detectable and at the end by the planet becoming undetectable.

Zone D is the line at the top of Zones A and B, where the observability period for the recovery of any planet is punctuated at the start by the stellar viewability window opening, and at the end by the stellar viewability window closing. Zone D includes the permanently detectable planets, in face-on orbits.

S*tarVault* contains these plots for all 136 stars on the target list.

4. SUMMARY, COMMENTS AND CONCLUSIONS

We have studied the requirements for planetary discovery using *TPF*. Although the focus has been on *TPF-C*, the general conclusions apply to *TPF-I*, as well. Our science-operational concerns do not apply to *TPF* observing planets found by other means—the *Space Interferometry Mission* (*SIM*) or ground-based radial-velocity searches, for example.

We have found ample indication that the science operations of *TPF-I* and *TPF-C* will face significant challenges eliminating background confusion and avoiding wasted observations in the attempt to recover previously discovered planets.

In the case of *TPF-C*, brightness will vary strongly in the detectable portion of the orbit for some planets. In principle, such brightness variations should be helpful for discriminating background sources and gaining orbital information useful for planetary recovery. We expect this phase effect to be more helpful for stars with better resolved habitable zones. Studies in progress should clarify and quantify the benefits. We expect to find a smaller, less helpful phase effect for *TPF-I*, at least for planets that efficiently redistribute the energy deposited by starlight, by rotation or atmospheric circulation.

Adequate astrometric accuracy addresses these challenges. Studies are underway to determine what astrometric accuracy is required to make science operations tolerably efficient and robust against the risks posed by undifferentiated confusion and guesses about recovery based on poor orbital information. These studies will use Monte Carlo methods to model aspects of the issues as well as the mission itself. The results will determine the minimum scale of the *TPF* mission.

Instrument designers are encouraged to study and report the expected astrometric performance of their *TPF-I* and *TPF-C* instrument concepts when measuring the star-planet separation and position angle.

Modeling similar to that shown here for the *discoverable* subset, could be performed on the *found* subset of planets—those consistent with the *measured* apparent separation and brightness of a particular discovered planet. This would more closely emulate the science-operational situation.

5. ACKNOWLEDGEMENTS

We salute Christian Lallo for his expertise in aiding the computations, graphics, and website. We thank Sharon Toolan with her help preparing the manuscript. JPL contract 1254081 with the Space Telescope Science Institute provided essential support.