Optimized Program Completeness

Sarah Hunyadi, Stuart Shaklan, Robert Brown

November 28, 2005

ABSTRACT

Overall program completeness is optimized by selecting the ordered set of observations and targets that maximizes the efficiency (completeness per time) for the ensemble of stars. We describe the optimization approach and report on completeness sensitivity to instrument throughput, inner working angle, instrument sensitivity, observational overhead, exo-zodiacal brightness, and target revisit constraints.

1. INTRODUCTION

Completeness was previously defined in Brown (2005) and Hunyadi et. al. (2005) as the fraction of potential planets in the habisphere that are detected. The cumulative completeness may then be defined as the sum of the completeness values integrated over all observations of the star. In this memo we will examine program completeness, which equals the sum of cumulative completeness over all stars. We define a method of observation by which stars are visited multiple times in order to maximize cumulative completeness.

Program completeness differs from single visit completeness (Brown, 2005 and Hunyadi et. al. 2005) in that there is potentially greater than one visit per star. In this memo we explore optimized program completeness and the effect of different stellar and instrument parameters on the final completeness results. Optimization serves to maximize the total number of detectable planets observed (cumulative completeness) in a limited integration time, with multiple visits to the most productive stars. The mission parameters are given in Table 1.

Symbol	Value		Quantity
D	8	m	long axis of the telescope mirror
d	3.5	m	short axis of the telescope mirror
IWA	65.5	mas	inner working angle with dither effect
Δmag	25		limiting delta magnitude sensitivity
λ	550	nm	central wavelength
$\Delta\lambda$	110	nm	bandpass
t_o	0.473		optical throughput
t_m	0.675		mask throughput
Ly^2	0.34		Lyot throughput
n_x	28.6		noise pixels
Ω_x	1.18E-15	steraidans	solid angle of criticaly sampleing pixels at central wavelength
ξ	0.001	$sec^{-1}pixel^{-1}$	dark count rate
ζ	5.00E-11		uniform contrast level in detection zone
φ	85	degrees	solar avoidance angle
R	2	$pixel^{-1}$	read noise
N_p	1.00E+04		number of planets populating the habisphere in random orbits
G	6.77E-11	$m^3kg^{-1}s^{-2}$	gravitational constant
M_{sun}	1.98E+30	kg	mass of the sun
r_e	4.85E-05	AU	radius of the earth
p_e	0.33		geometric albedo of the earth
ϖ	0.3333		fraction of total mission time used as integration time

Table 1. Instrument and mission parameters for the TPF-C mission.

2. COMPLETENESS

One of the main differences between program completeness and single visit completeness is that we now account for the progression of time. Program completeness must take into account two temporal factors:

a) The telescope orbits around the sun, restricting the observation of candidate stars to when they are visible to the telescope.
b) The planets populating the stellar habispheres orbit around their parent star.

For single visit completeness, we assumed that the stars were always visible because we were only taking a snapshot of the habisphere at a given time over a year and within that year each star would be visible at least once. This is no longer a valid assumption as we are now accounting for the orbit of the Earth and observing over a 3 year mission period. As the Earth and telescope orbit the sun (the baseline telescope position is at L2), each candidate star is visible for a period of approximately six months. This availability time is slightly reduced due to a solar avoidance angle restriction that limits viewing to an 85 degree cone in the anti-sun direction. This observation window shifts in space as mission time advances, allowing new stars to enter the window and some previously available stars to exit the window of availability.

Although different in effect, the change in planet visibility also occurs over a window of time. The observability of a planet in a non-face on orbit depends on its motion. While a

face-on orbit planet has the same brightness regardless of its phase, the brightness of non-face on orbit planets changes periodically. As the planets rotate behind the star they become brighter (position 1 in Fig. 1) and as they move around the front of the star they become dimmer (position 2 in Fig. 1). Therefore a planet behind the star can be observed at a lower sensitivity limit that those in front of the star. This change in magnitude factors into the optimization of the visit timing for a particular star.

Fig. 1. Diagram showing the effects of planet rotation on brightness. This is the case for a non-face on orbit. Planets become brighter as they move behind the star (position 1) and dimmer when they move in front of the star (position 2).

Much like single visit completeness (Hunyadi et. al., 2005), the habisphere is again modeled with 10,000 planets (N_p) in random and evenly distributed orbits with uniformly distributed eccentricities from (0, 0.35). The planets have a position given by the three Euler angles that define the planet location (Brown, 2004), ϕ (rotation about the line of sight, $0 \le \phi \le 2\pi$), θ (rotation about the plane of the sky, $0 \le \theta \le 2\pi$), and ψ (second rotation about the line of sight, $0 \le \psi \le \pi$). In single visit completeness this set of parameters completely defined the planet position. Implementing a program completeness approach, the Euler-defined planet location (described in the single visit memo, Hunyadi et. al. 2005), is taken to be the initial position of the planets, as we now account for their orbital motion. They have a period given by:

$$T^2 = \frac{4\pi^2 a^3}{GM_{star}} \qquad (1)$$

where a is the semi-major axis of the orbit, G is the gravitational constant and M_{star} is the mass of the star in solar masses. The mean anomaly (or phase) is now time dependent and given by:

$$M = M_0 + \frac{2\pi t_i}{T}, \qquad (2)$$

where t_i is the timing of the visit and M_0 is the initial phase of the planet that is random on $[0, 2\pi]$. This factor can then be utilized to determine the true anomaly of the planet

and therefore the initial phase (for more details see Hunyadi et. al., 2005.) By incorporating the time dependence in eq. (2), the phase of the orbit, the position (projected distance) and the delta magnitude (Δmag) change with time. The equations for the Δmag and the projected distance from the star (r_p) are given by (Hunyadi et. al., 2005):

$$\Delta mag = -2.5\log(r_e^2) - 2.5\log(p_e \Phi_l(\beta)) + 5\log r\,, \tag{3}$$

$$r_p = \sqrt{x_{fzx}^2 + y_{fzx}^2}\,, \tag{4}$$

where r_e is the radius of the earth, p_e is the albedo, $\Phi_l(\beta)$ is the planetary illumination, r is un-projected distance from the planet to its parent star in the planetary coordinate system, x_{fzx} is the rotated x coordinate and y_{fzx} is the rotated y coordinate.

A plot of projected planet distance vs. Δmag may be seen in Fig. 2. This figure was first presented by Brown (2005) and is recreated for the orrery of N_p planets.

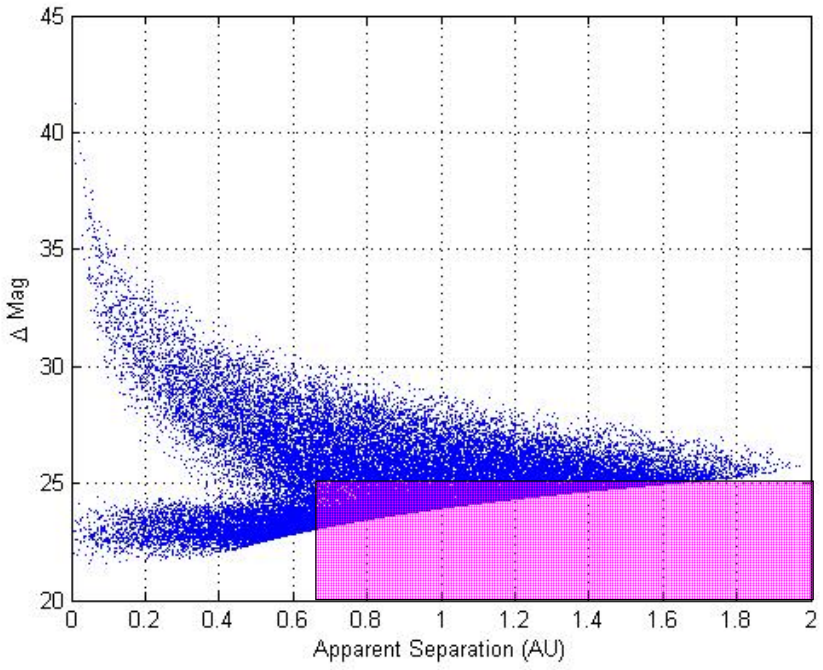

Fig. 2. The planetary probability density distribution for earth like planets on habitable orbits for a solar type star. The edges of the pink box indicate the current sensitivity requirements for the TPF-C mission.

For program completeness, the planetary probability distribution for earth-like planets in habitable orbits is initially identical to that for the single visit completeness case. For lower luminosity stars, the shape of the distribution in Fig. 2 stays the same, but moves down in the ordinate direction and is compressed in the abscissa towards lower apparent separation (projected distance). For high luminosity stars, the distribution shifts up in Δmag and extends further out in projected distance. The TPF-C mission has a maximum Δmag and a minimum inner working angle set by the telescope dimensions and stability parameters. These parameters define a box (Fig. 2), the height of which is the Δmag sensitivity limit. The baseline Δmag sensitivity limit is $\Delta mag = 25$.

The left side of the pink box is the IWA, which is currently set to $4\lambda/D$, giving an IWA = 57.0×10^{-3} arcsec. However, the IWA is effectively IWA = 65.5×10^{-3} arcsec after accounting for the shape of the mask and telescope rotations. The projected inner working angle or minimum apparent separation (the x-axis in Fig. 2) is defined as the IWA multiplied by the distance to the star in parsecs.

The limiting Δmag and projected IWA define a box, as shown in Fig. 2. The number of planets (N_b) that fall within this box divided by the total number of planets (N_p) is the single visit completeness:

$$Comp_{SV} = \frac{N_b}{N_p}. \tag{5}$$

In our previous memo (Hunyadi et. al. 2005) we showed curves for completeness vs. integration time for a single visit. However, these curves are only valid for the first visit to the star. For subsequent visits, completeness is determined by the detectable planets that were not previously viewed. Thus the completeness equation per visit becomes:

$$Comp_i = \frac{N_{bi} - N_{b0}}{N_p}, \tag{6}$$

where N_{b0} is the number of planets that are visible that were already observed and N_{bi} is the number of planets that are visible at a given time, t_i. Therefore, for each visit to the star, completeness vs. Δmag must be re-calculated based on the number of planets that have not yet been observed.

3. OPTIMIZATION

For program completeness, the visit quantity and quality (sensitivity) significantly impacts cumulative completeness for an ensemble of stars. Optimization serves to maximize cumulative completeness by trading stellar integration time and the number of visits. In order to efficiently solve this maximization problem the following assumptions were made:

1) Integration times are discretized into 1-hr blocks
2) The set of observations is discretized into week periods called visits.
3) The planets and Earth are assumed to be static for the duration of a visit.

A series of visit periods totaling one year of integration time are spread over a three year observing program. For a given visit spacing, or minimum revisit time, a fraction of the total mission time is used for integration as:

$$IT_i = \varpi(VS), \tag{7}$$

where IT_i is the integration time per visit period in weeks, VS is the visit spacing in weeks and ϖ is the fractional proportion of the visit spacing, which is set to be 1/3 of the visit spacing for this optimization. For example, a visit spacing of 3 weeks would lead to one week of integration time for every three week period. The effects of the different visit spacing on cumulative completeness are shown and discussed in Chapter 4.

3.1 Two-Visit Completeness

We will initially examine the case of two visits for a given star, which was first presented by Brown (2004). Figure 3 is a plot of the two visit completeness vs. stellar visit spacing. The three curves in Fig. 3 represent different luminosity stars with the same obscuration ($a_0/\sqrt{L}$). The curves in Fig. 3 are oscillatory in nature, with different periods of oscillation depending on the luminosity. For a solar type star (L=1), the first and second valleys in the curve occur when planets rotate 180 degrees from their original position (first valley at the point marked in red at 6 months) and when they return to their original position, 180 degrees later (second valley at the point marked in red at 1 year). The peak visit completeness occurs at a revisit time of roughly 3 months (shown in black in Fig. 2). For different luminosity stars, the peak two visit completeness occurs at different times. For instance, a low luminosity star (L = 0.1) has a peak revisit time of 2 weeks while a high luminosity star (L = 4) has a peak revisit time of 34 weeks (also marked in black in Fig. 3).

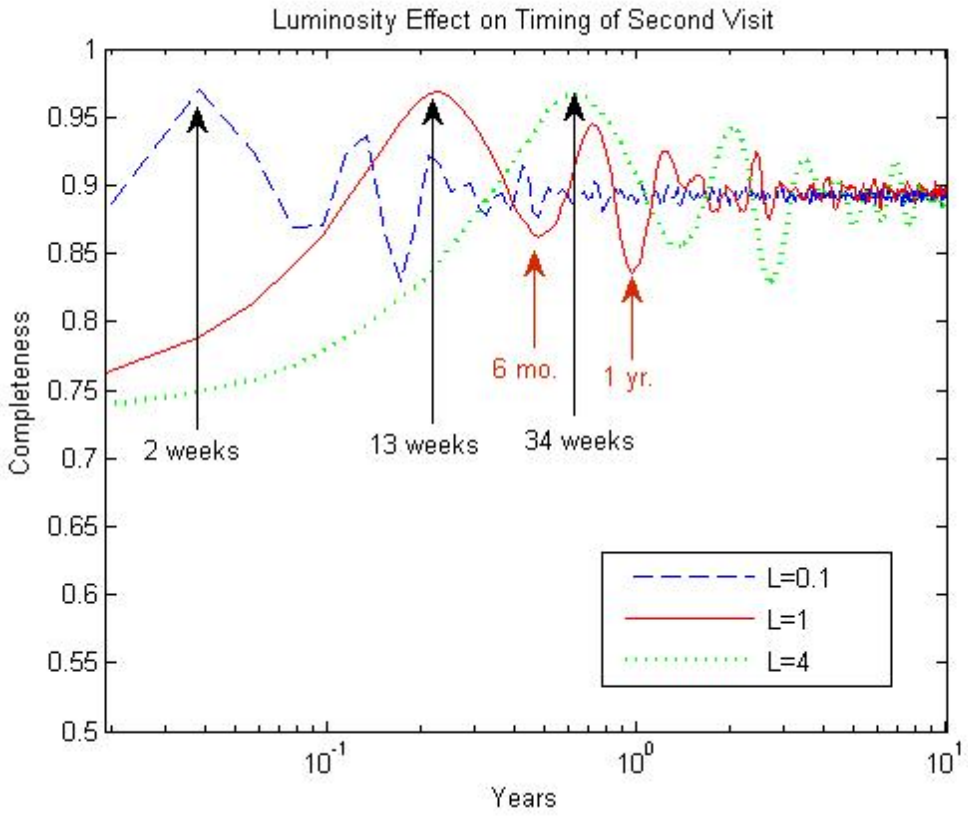

Fig. 3. Second visit completeness differences for stars of different luminosities with the same obscuration.

Because the curves are oscillatory in nature, the discretization and freezing of planets adds an error into the simulation when the orbital period of the planet is not adequately sampled. This has the effect of smoothing out the peaks and valleys of the curves in Fig. 3 and has a larger effect for low luminosity stars, as the oscillations occur on a much shorter time scale than for higher luminosity stars. In this word, we find many cases of multiple revisits and revisits spaced by times grater than the optimum. Thus, the smoothing effect is of little consequence.

3.2 Program Completeness

This oscillatory behavior also occurs for more than two stellar visits. For larger visit number, the visit curves have higher level structure, but remain oscillatory in nature. Therefore, smoothing of the curve exists for three or more visits with the current assumptions, as in two visit completeness.

Since completeness oscillates with visit timing for any number of visits, an optimal revisit time exists for every star at the peaks of the individual stellar completeness curves. One optimization method involves maximizing completeness on a star by star basis. While this optimizes the completeness of a given star for a given number of visits, it does

not necessarily optimize the program completeness because it does not optimize over the collection of stars relative to each other.

Given a constraint on total observing time, it is also possible to optimize over a collection of stars by determining the length of observation and set of stars that offers the highest average rate of planet detection in each visit period. Therefore, we optimize visit by visit. We then adjust the minimum revisit time to search for an optimum. Figure 4 is a schematic of our approach, described in the following paragraph.

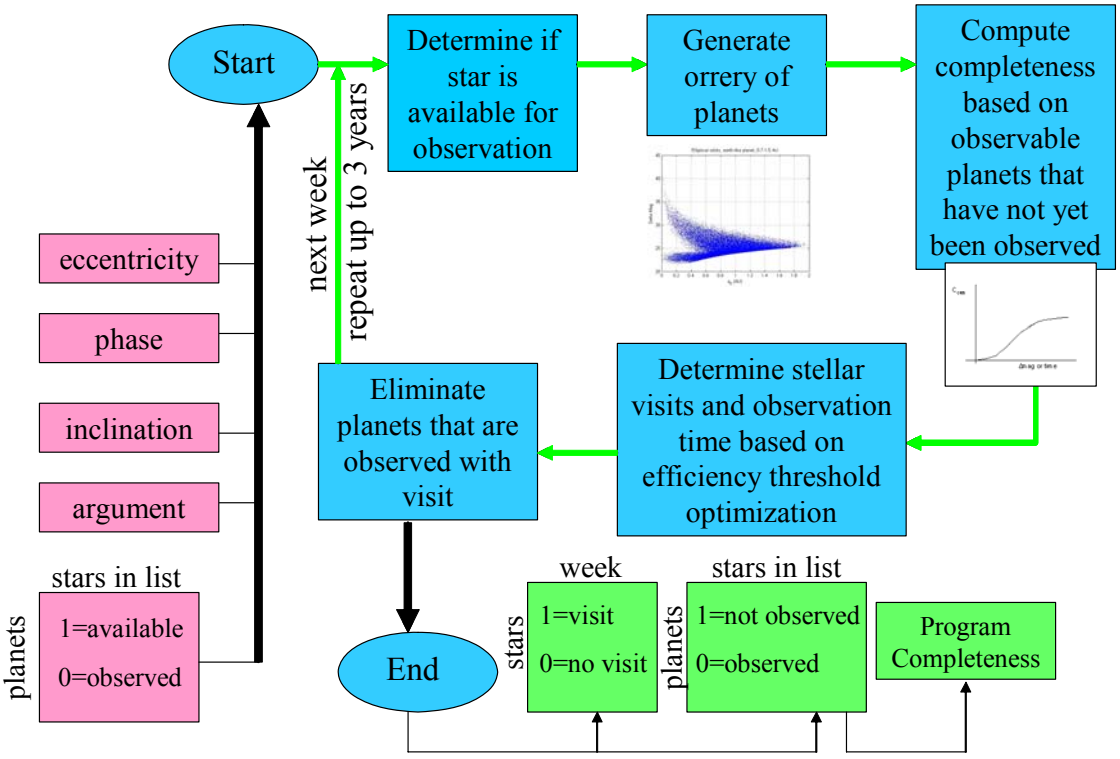

Fig. 4. Program completeness diagram. Pink boxes indicate inputs and green boxes indicate outputs.

At the start of the optimization process, we set the input parameters that define the planetary orbits of the orrery surrounding each star (depicted as pink boxes in Fig. 4). Next, a determination is made on the availability of a star, depending on its position in the ecliptic. If the star is available, an orrery of N_p planets is generated and a curve of completeness vs. time is calculated based on the planets that have not yet been viewed. Next, an efficiency threshold cutoff (ETC) optimization is performed on the stars that are available (Hunyadi et. al., 2005). This optimization process sets an ETC that evaluates the instantaneous efficiency (change in completeness per hour) of a star and eliminates all time that falls below the ETC. The stars are then sorted by average efficiency (completeness per total integration time remaining) and those stars whose integration times sum to less than the integration time limit (IT_i in eq. 7) are selected for a visit. The "detected" planets are then eliminated from the list of possible planets for the visited stars and we return to the beginning of the optimization process (labeled with a Start ellipse in Fig. 4) for subsequent visit periods.

In this first iteration of the program completeness loop, all planets around all stars are available. For the second visit period, this is no longer the case. The only planets available for observation are:

1) planets that were not previously visible that rotate into the observable habisphere
2) planets that were not observed in the first observation because the star was observed to a lower Δmag limit (high ETC eliminates time and potential observable planets), and are now visible due to a smaller ETC (lower acceptable efficiency).
3) planets around stars that were not visible due to the solar avoidance angle.

Completeness curves and the ETC are re-calculated based on the remaining observable planets. The optimal value of the ETC depends on both the collection of stars that are selected for observation and the observational history of a given star. For the first several iterations, the ETC is large as there are many stars with high initial efficiency that are competing to make the cut. As time progresses, the ETC becomes smaller as stars become less productive (lower efficiency). This implies that stars that have already been visited have much lower maximum completeness per visit than in the first iteration and, therefore, have lower incremental and average efficiency. Therefore, only stars that are bright and nearby (with short integration times) or low luminosity (with many new planets that rotate into the observable habisphere rapidly) will be visited repeatedly. In each visit the observed planets are eliminated for the stars that are observed by ETC optimization.

This process (shown with green arrows in Fig. 4) is repeated every *VS* weeks over a three year mission. At the end of this three year period, we exit the loop and obtain two outputs (shown in green boxes in Fig. 4). The first output delineates the weeks in which the star was visited. The second output is a matrix listing which planets, around which stars, were detected over the course of the mission. By summing the columns of this second output, the completeness per star may be obtained. By then summing this result, we may obtain the program completeness (third green box in Fig. 4).

4. RESULTS

4.1 Baseline Results

The columns in the output of program completeness in Fig. 4 represent the cumulative completeness gain over time ($Comp_i$) for a given visit spacing and can be shown pictorially (as in Fig. 5). According to eq. 7 with $\varpi = \frac{1}{3}$, visit spacing less that 3 weeks would lead to less than one week of integration time (IT_i) per visit period. Some visited stars require up to this limiting IT_i and are removed from the candidate list for smaller IT_i. Therefore, a baseline visit spacing of 3 weeks was chosen to allow for longer integration time stars to make the list of possible candidates while minimizing the freezing error. A depiction of the 3 week baseline visit program is shown in Fig. 5.

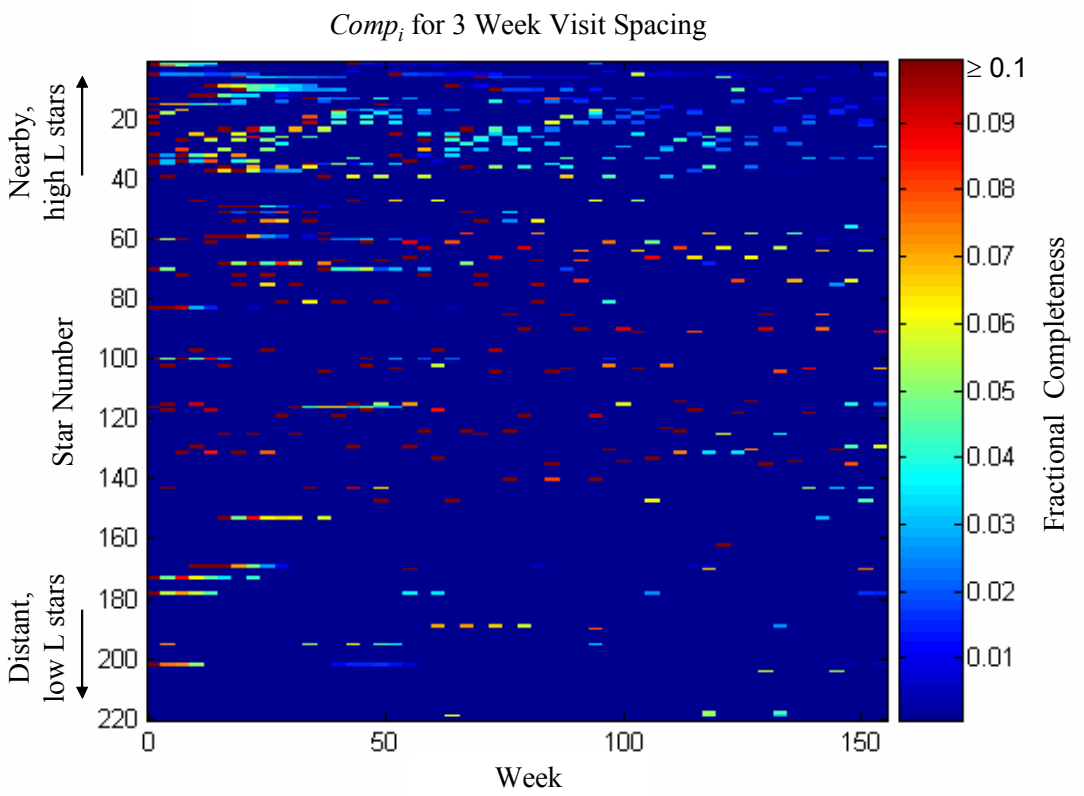

Fig. 5. $Comp_i$ for 3 week visit spacing. Deep red indicates a $Comp_i$ of greater than or equal to 10% per visit. Blue indicates a visit with between 0 - 4% $Comp_i$ per visit.

Each colored bar in Fig. 5 represents a visit to a star after optimization. The columns represent visit periods, each 3 weeks long. The rows are star numbers, sorted by $d/\sqrt{L}$ (angle of HZ) so low luminosity, distant stars are at the bottom of the chart. The color of the bars depicted in Fig. 5 represents the $Comp_i$, (as in eq. 6) or number of planets out of the sample of 10,000 that are observed per visit. Dark red indicates an observation with large completeness gain (large number of planets observed) while blues indicate a visit of little completeness gain (few new planets observed). In Fig. 5, large completeness gain is defined by an observation when 1,000 or more planets (fractional completeness = 0.1) were observed in a single visit. Some stars are visited in consecutive multiple visit periods, shown by several bars of different color linked together. Stars are visited only when their efficiency is high when compared to other possible candidate stars. This means that stars that are visited multiple times are those that can give high completeness gain ($Comp_i$) in a short time.

The fact that some stars are visited more than ten times, and further, that they are visited in consecutive visit periods is initially counterintuitive. It would seem to be more beneficial to spend time visiting a new unobserved star instead of revisiting a star that has already been observed. However, this is not the result obtained in this simulation. There are two types of stars that contribute to this unexpected behavior. The stars in the upper left of Fig. 5 are revisited multiple times because they are nearby, well resolved stars requiring short integration times. These stars require short integration time (< 10 hours)

to reach the Δmag limit quickly, irrespective of the luminosity (Fig. 6b). This is true even as the subsequent visits become less productive (lower $Comp_i$ depicted by visit lines that change from red to blue in color as we move from left to right in Fig. 5). Stars further down in the chart are also visited multiple times, but for a different reason, which is best understood by observing the shapes of two completeness curves for two different luminosity stars (Fig. 6).

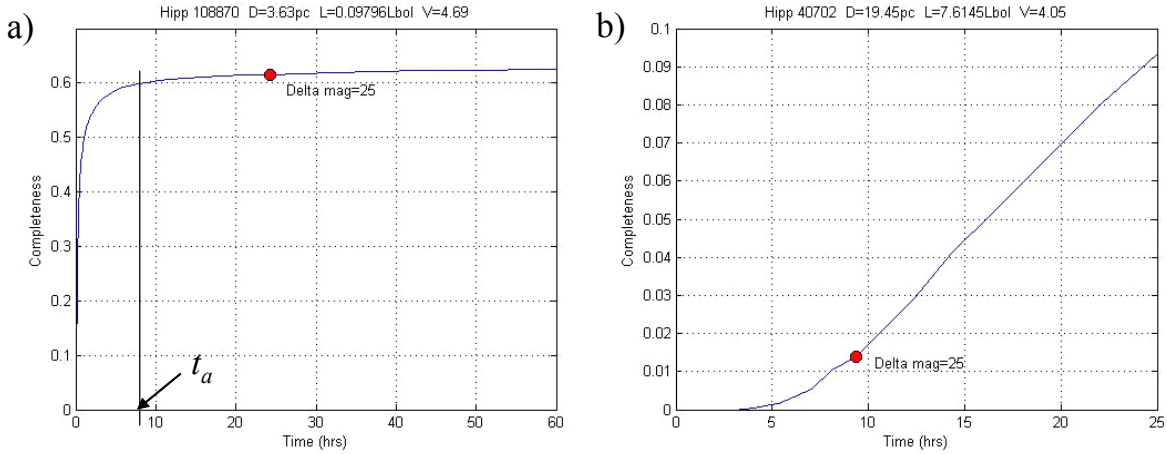

Fig. 6. Plots of completeness vs. integration time. a) Completeness curve for a lower luminosity star. The time beyond which there is minimal completeness gain is given by t_a. b) Completeness curve for a high luminosity star. The red points indicate the point of limiting Δmag.

The curve in Fig. 6b continues to rise after the limiting Δmag cutoff, indicating that there is significant completeness gain available from observing to a deeper sensitivity limit. On the other hand, the completeness curve in Fig. 6a rises rapidly and then levels off, giving little completeness gain for times longer than time t_a. Therefore, we may cut integration time from low luminosity stars without loosing significant $Comp_i$. In addition, the habitable zone of low luminosity stars lies closer to the star than for high luminosity stars. This means that the planets orbit the star more rapidly and have shorter period orbits. The rate of new planets entering the observable habisphere is therefore greater for these low luminosity stars. This rapid regeneration of possible targets combined with the ability to cut integration time without loosing significant completeness allows us to view these stars multiple times at integration time t_a rather than visiting other stars that require longer integration times. This is depicted in Fig. 7, as cutting integration time is equivalent to observing to a lower Δmag.

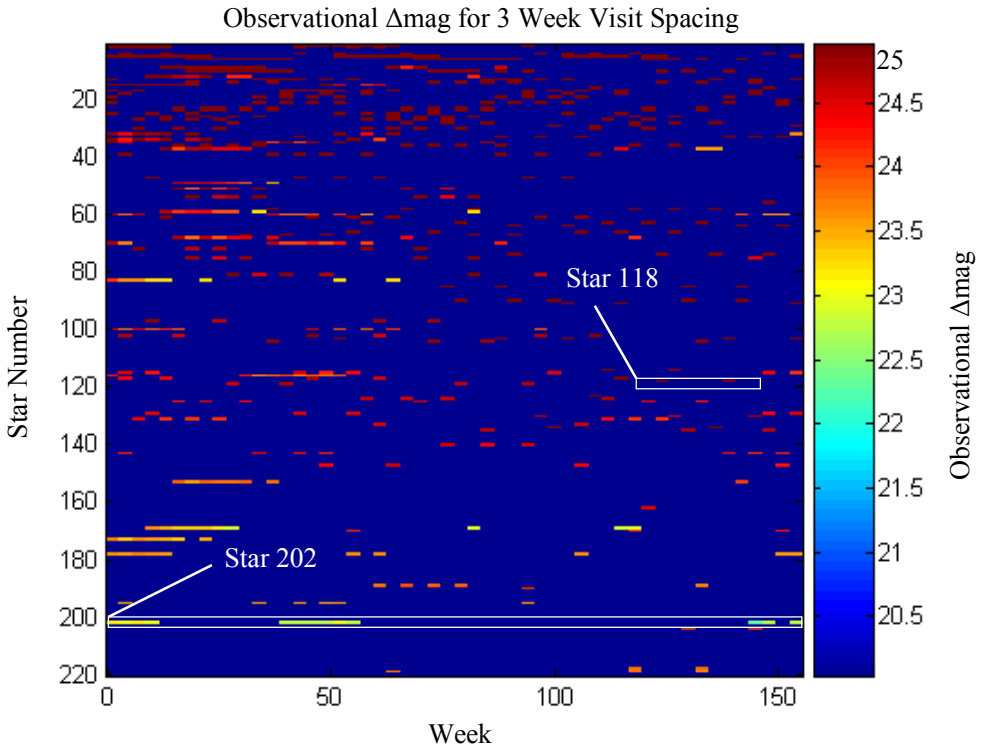

Fig. 7. Limiting Δmag associated with each stellar visit

As in Fig. 5, the colored bars in Fig. 7 represent observations. Darker red lines indicate an observation up to or close to the limiting Δmag. The top left of the chart shows many observations to deeper magnitude as the integration times are shorter for these nearby, bright stars. Thus, they have higher efficiency at deeper Δmag and can be observed deeper at little mission time cost. The lighter (green to orange) lines indicate an observation at a lower sensitivity (lower Δmag). Therefore, lighter lines also indicate a star from which a large portion of time has been cut relative to t_{max} (the time to reach the Δmag limit). One such low luminosity star is Hipparcos# 108870 (star number 202). The observations for this star are at most up to Δmag = 23.4, indicating that even in the first observation (when the maximum completeness and therefore incremental efficiency is highest) removing integration time was beneficial. This star was also observed in sequential observation periods as shown by the long yellow lines at row 202 in Fig. 7. This indicates significant (when compared to other stars) completeness gain every time it was visited.

Towards the end of the mission (week > 100) new stars that were previously unobserved enter the visit plan. This is because the observable planets for the early or good stars have been eliminated from the list and there is minimal gain from revisits. At the end of the mission the stars requiring longer integration time are now above the ETC and warrant an observation. One such star is Hipparcos# 32439 (star number 118). It has a luminosity of 1.73 L_{bol} and a distance of 17.85pc from the Sun, leading to long integration times and a completeness curve that continues to slowly rise beyond Δmag=25. It is visited only twice, in visit week 124 and visit week 139.

A scatter plot of the number of visits per star is shown in Fig. 8 and a histogram of the number of visits is shown in Fig. 9.

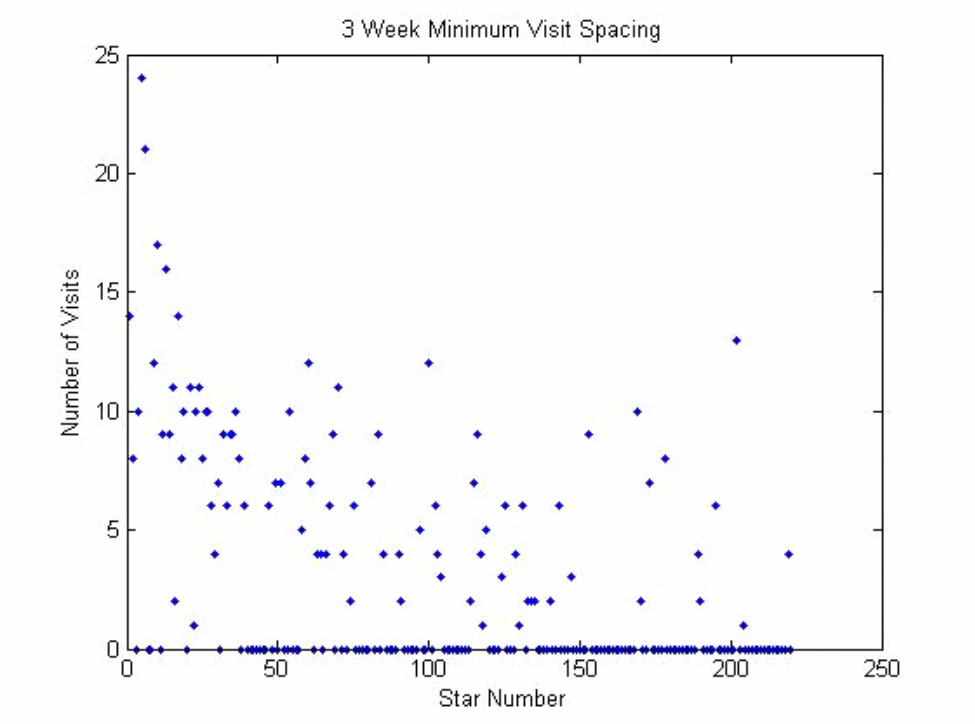

Fig. 8. Scatter plot of the number of visits per target.

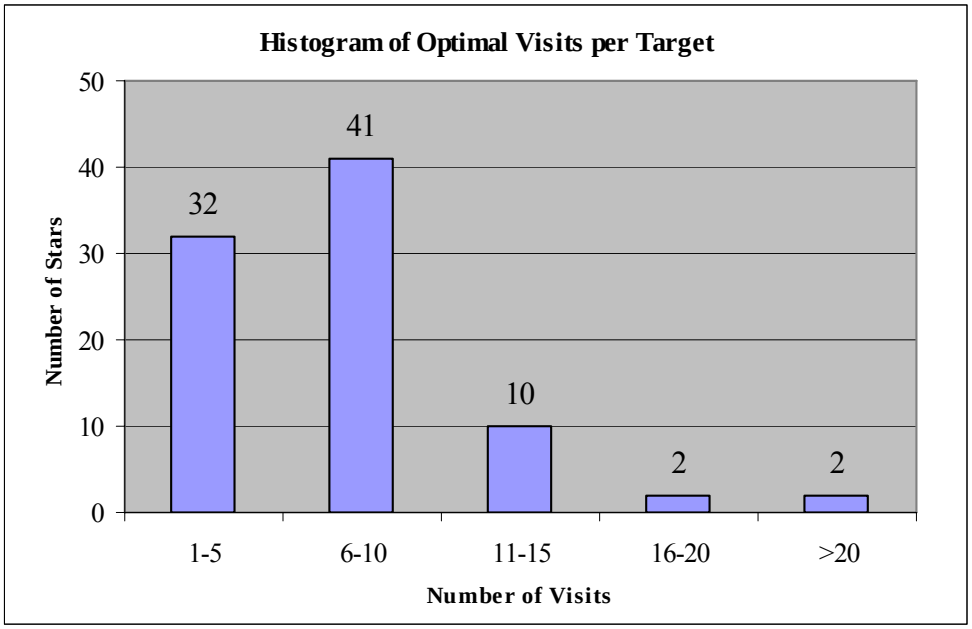

Fig. 9. Histogram of the number of visits per target, given by the baseline mission parameters and visit spacing.

Only 87 stars out of the top 220 stars (sorted by $\sqrt{L}/d$) are visited one or more times. The stars that are not visited are shown as points along the x-axis in Fig. 8. The maximum number of possible revisits in the baseline 3-week visit spacing scenario is 25. One star (Hipparcos# 2021) is visited almost every time it is available, with 24 total visits. There is a broad range of visit number, shown best in Fig. 9. The majority of stars are visited 6 or more times. Since previous program completeness estimations limited the number of visits to six or less (Shaklan and Lisman, 2004), the large number of visits per star is unexpected and is the defining result of this iteration of optimizing program completeness.

The significant change in visit strategy (observing one star many times instead of many stars fewer times) leads to a significant increase in the cumulative completeness when compared to the single visit case. For optimized single visit completeness at Δmag = 25, the cumulative completeness was ~ 26 HZs searched (Hunyadi et. al., 2005), while the cumulative completeness for the baseline optimized program completeness is greater than 46 HZs searched (both assume one-year total integration time).

4.2 Visit Timing

The optimized program completeness may change with different choices of visit strategy. In this memo, visit strategy may be defined as the collection of stellar visits resulting from the program completeness optimization. The strategy is therefore dependent on visit spacing (spacing of observation periods), minimum revisit time (time between allowable star visits) and stellar efficiency. In the baseline case, the minimum revisit time and visit spacing are set to be the same, as every time a star is available for a visit, a visit is permitted. It is possible to restrict the availability of a star to certain visit periods, thereby limiting the number of times that it is visited (e.g. if a star is visited in week 1, the next possible visit occurs is week 10 for 9 week minimum revisit spacing). This restriction can also be applied to a collection of stars. For every star a visit is allowed only after a prescribed amount of time has passed since the previous visit. This forces the optimization to select a different group of stars for observation that were not observed in the previous visit period, leading to a new visit strategy. Figure 10 shows a plot of cumulative completeness over time for different choices of minimum revisit time.

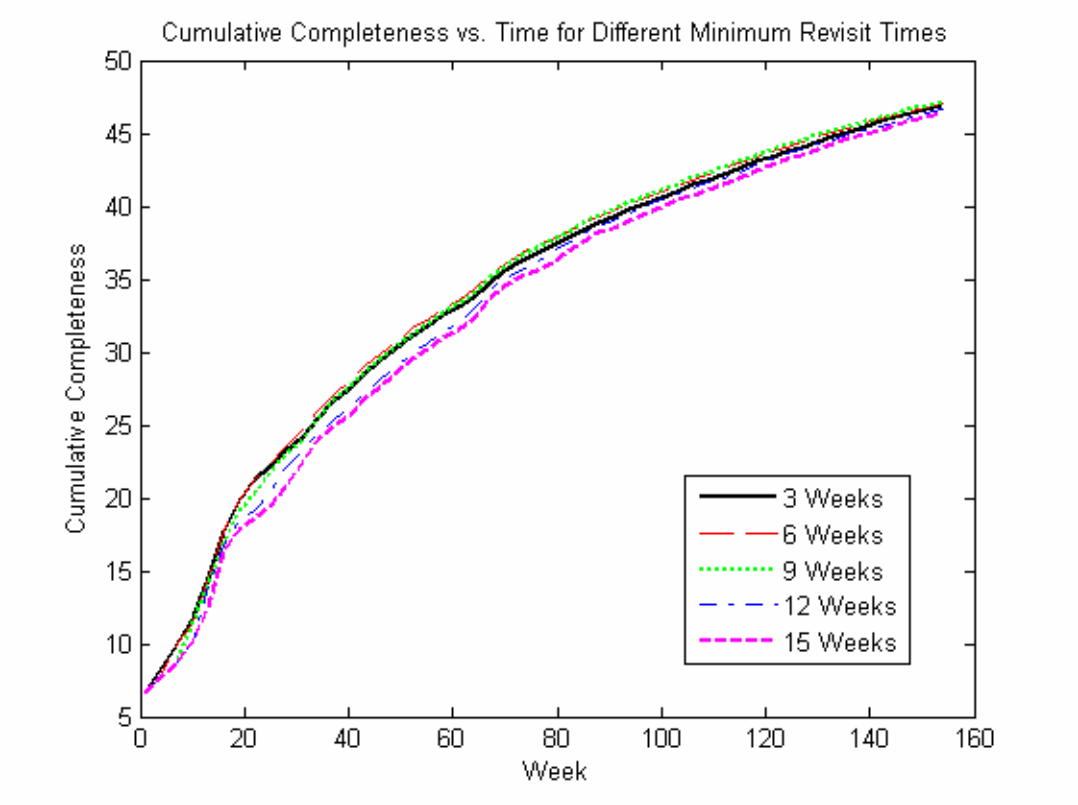

Fig. 10. Plot of total completeness over time for different minimum revisit times. The baseline case is 3 weeks (solid black curve), where the revisit time and visit spacing are the same.

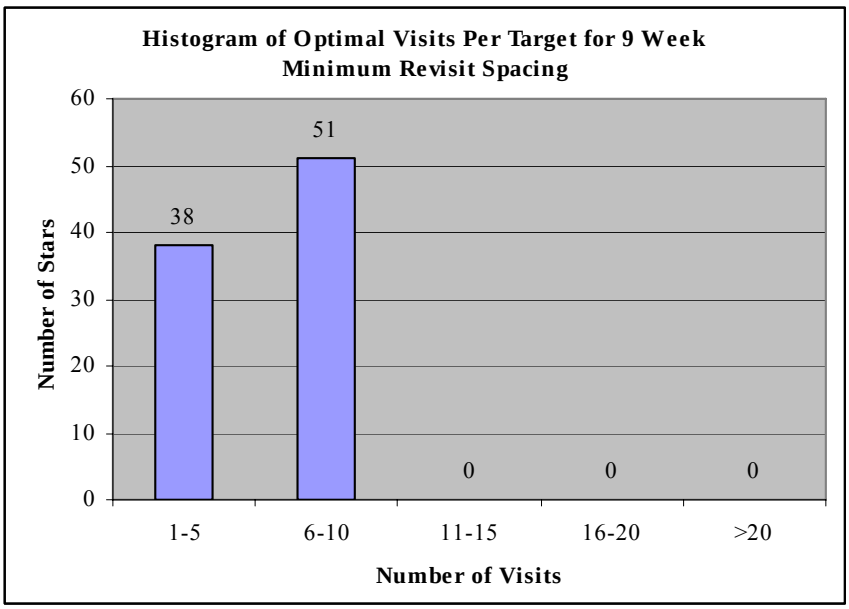

Fig. 11. Histogram of the number of visits per target, given by the baseline mission parameters and 9 week minimum revisit spacing.

Minimum revisit time was varied from being the same as the baseline visit spacing (3 weeks) to 15 weeks between visits. The curves in Fig. 10 show that even for large minimum revisit time (12 – 15 weeks) there is little effect on cumulative completeness, especially at the end of the mission. There is less than 0.3 searched HZ difference between the baseline spacing (shown in black in Fig. 10) and the highest cumulative completeness value (shown in green in Fig. 10 for 9 week minimum revisit time). The list of visited stars also does not change significantly with minimum revisit time (at 15 weeks there are 5 additional stars visited that were not visited at 3 weeks). For large revisit spacing the individual visits are more productive, but there are fewer total visits so the cumulative completeness remains the same.

The optimized cumulative completeness may also change with choices of different visit spacing. Visit spacing is varied from 2 weeks to 7 weeks, where the baseline occurs at 3 weeks. A plot of the number of habitable zones (HZs) searched vs. visit spacing is depicted in Fig. 12.

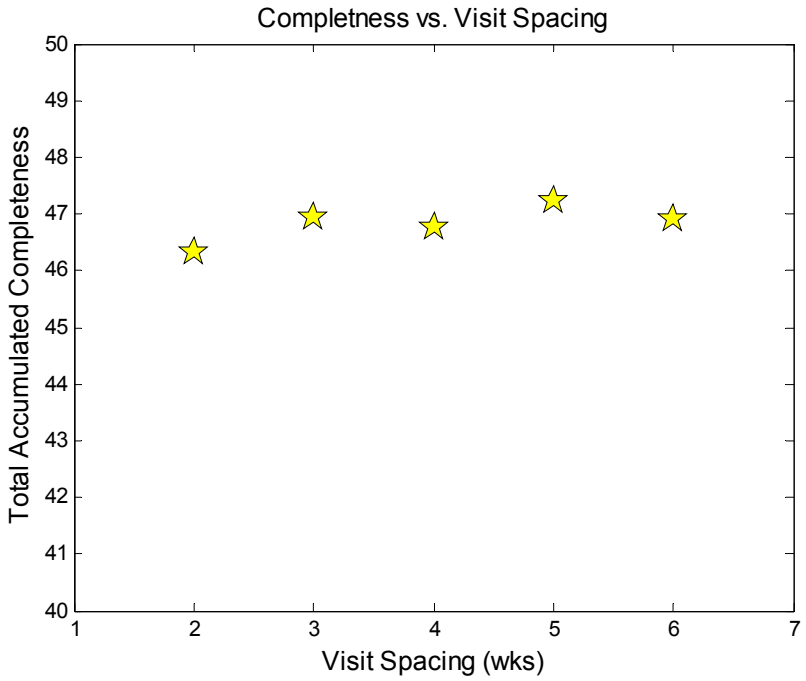

Fig 12. Comparison of completeness results for different visit spacing.

In two-visit completeness (as presented in Chapter 3.1), the effect of visit spacing was very pronounced. Figure 12 shows the effect of visit spacing on program completeness for the optimized case. This effect contrasts that of the two-visit completeness in that there is minimal effect from changing the visit spacing. This is partly due to the freezing error which smoothes out the peaks and valleys of the visit spacing completeness curves. It is also due to the fact that different luminosity stars have different peak revisit times and the optimization occurs over the collection of candidate stars.

Although the cumulative completeness is not significantly affected by minimum revisit time or different visit spacing, it is affected by changing the mission parameters.

Exploring changes in these parameters allows us to determine the impact of future architecture choices.

4.3 System Requirements Trades

There are three types of parameters that affect completeness and architecture choices for the mission. These parameters are: strategic factors, tactical factors and system requirements. Strategic factors are factors that could impact the design of the mission if the factors were known in advance. These include zodiacal dust density, presence of known planets, and confusion sources. Tactical factors are factors that are not knowable in advance and affect the mission as it progresses. These include planetary detections, time required to confirm a planet detection and time required to characterize a planet after it has been detected (Brown, May 2005). System requirements define the constraints on the science mission and are determined by the mission objectives and engineering constraints.

By varying the system parameters, the range of scientific benefit of different architectures and systems may be explored. In this section we address the impact of system parameters. The explored system parameters include:

1) Throughput – the amount of light that is received at the detector relative to the total quantity incident on the primary mirror
2) Inner working angle – (IWA) – the smallest observable angle between the star and the planet
3) Overhead – the time spent maneuvering the telescope before observing
4) Limiting Δmag – the limit to the faintest planets observable, relative to the host star as determined by the stability requirements on the telescope.

The only strategic factor considered in this memo is exo-zodi level or brightness of the dust around the observed stars. However, since exo-zodiacal brightness and distribution are not currently well understood, these results will only indicate how the mission design is able to deal with different quantities of exo-zodiacal dust. A plot of the effect of exo-zodi brightness on cumulative completeness is shown in Fig. 13.

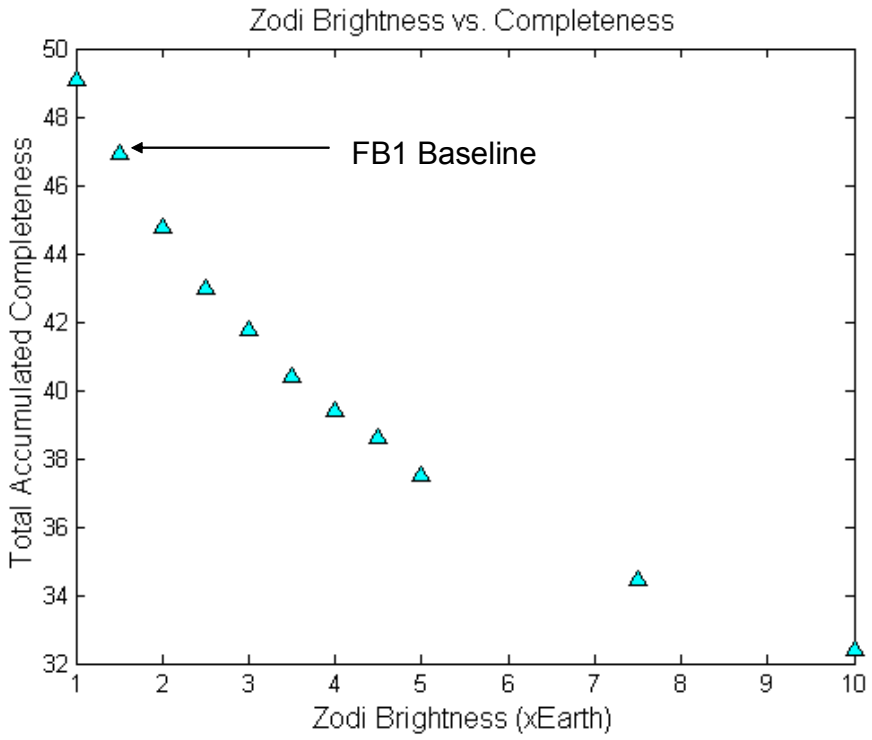

Fig. 13. Effect of increasing exo-zodi on cumulative completeness. The baseline is set to 1.5 times that of Earth zodi brightness.

The current baseline for exo-zodi is set to 1.5 times that of earth zodi, meaning that we see three times as much exo-zodiacal dust as Earth zodiacal dust (we look through the exo-zodi but only look out of our zodi). Earth zodi brightness is assumed to be 23 mag/arcsec^2. With an increase in the level of exo-zodiacal dust brightness, the integration time increases since more time must be spent per star to achieve the same SNR. This decreases the cumulative completeness as the zodi brightness increases. As shown in Fig. 13, there is a 20% loss of completeness at zodi density greater than 5x that of Earth zodi. If the requirement decreases from 1.5x Earth zodi brightness to 1x Earth zodi brightness, the cumulative completeness increases by 7%. Therefore, large changes in zodi brightness significantly impact completeness (20% change) and it is a strong driver at levels of 5x Earth zodi.

Noteworthy scientific gain (as opposed to loss as with zodi-brightness) may be obtained by increasing the system throughput. System throughput depends on the number of components in the coronagraph and the efficiency of each component. The Lyot stop diameter (which affects throughput and sharpness) was kept constant in the simulations.

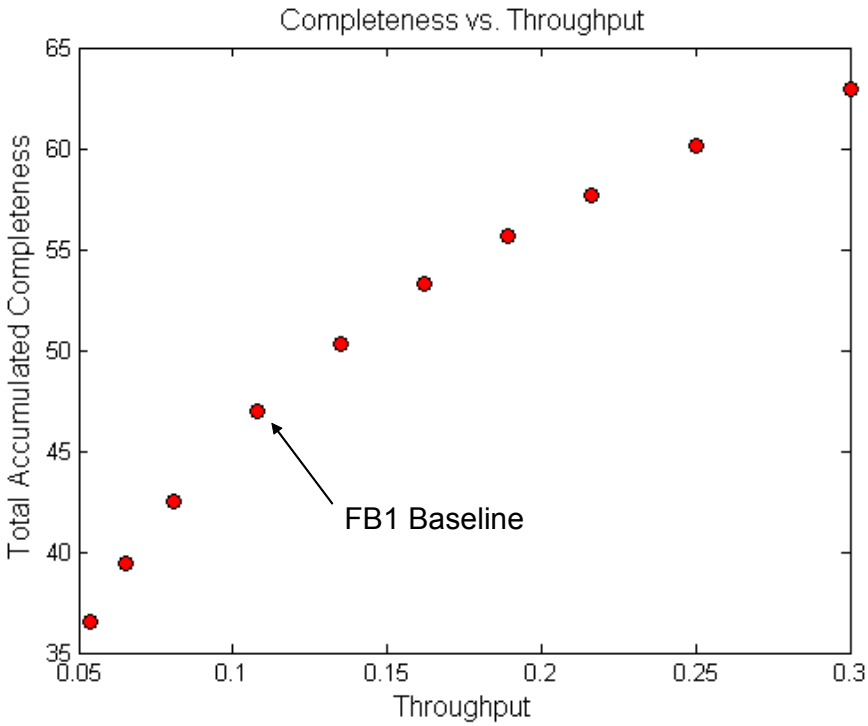

Fig. 14. Completeness dependence on system throughput. The baseline throughput is 10.8% (not including CCD QE).

Figure 14 is a plot of the number of HZs searched vs. system throughput. The baseline system throughput is 10.8%. An increase in throughput to 20.4% yields approximately a 20% increase in cumulative completeness. A decrease in throughput to 5.8% yields a 20% decrease in cumulative completeness. Alternative coronagraph designs that increase system throughput are currently under consideration.

There are also alternative designs that would decrease the inner working angle, which is another factor that may be studied in the architecture performance evaluations. The inner working angle is defined by the coronagraph architecture and instrument susceptibility to low order deformations. Figure 15 is a plot of the number of HZs searched vs. IWA.

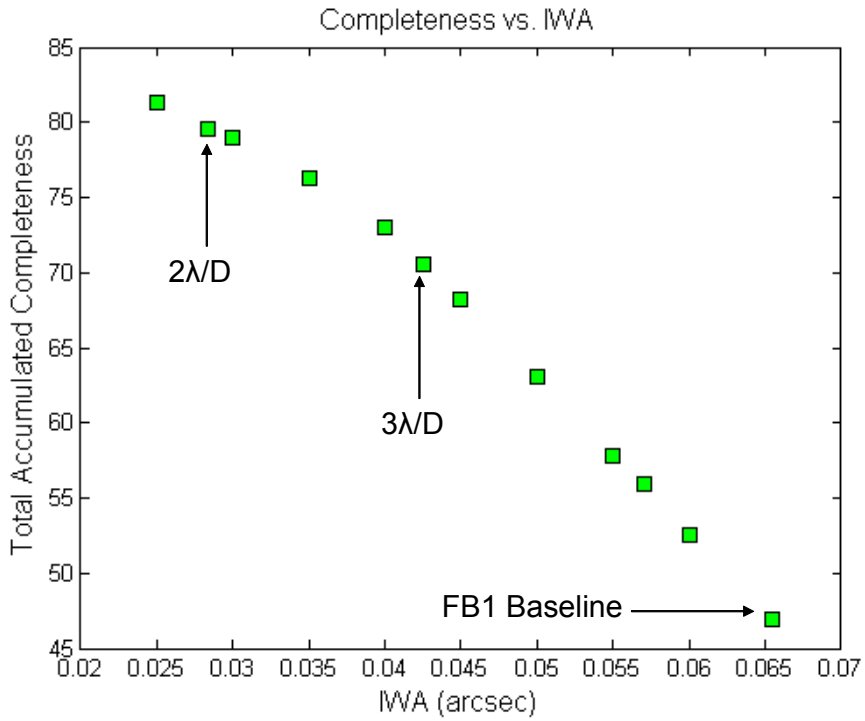

Fig. 15. Change in completeness from changing the inner working angle. The baseline is set to 65.5 mas. The labeled points assume λ=550nm and D=8m.

Decreasing the IWA leads to a rapid and significant increase in cumulative completeness. There is a 20% improvement in completeness by decreasing the IWA from 65.5 mas to 55 mas. By further decreasing the inner working angle to 3λ/d (42.5 mas) or even further to 2λ/d (28.4 mas), the cumulative completeness can be increased by 34% and 41%, respectively.

One additional factor that has a significant impact on system requirements, but does not strongly influence the science return is the limiting Δmag. A plot of completeness vs. limiting Δmag may be seen in Fig. 16.

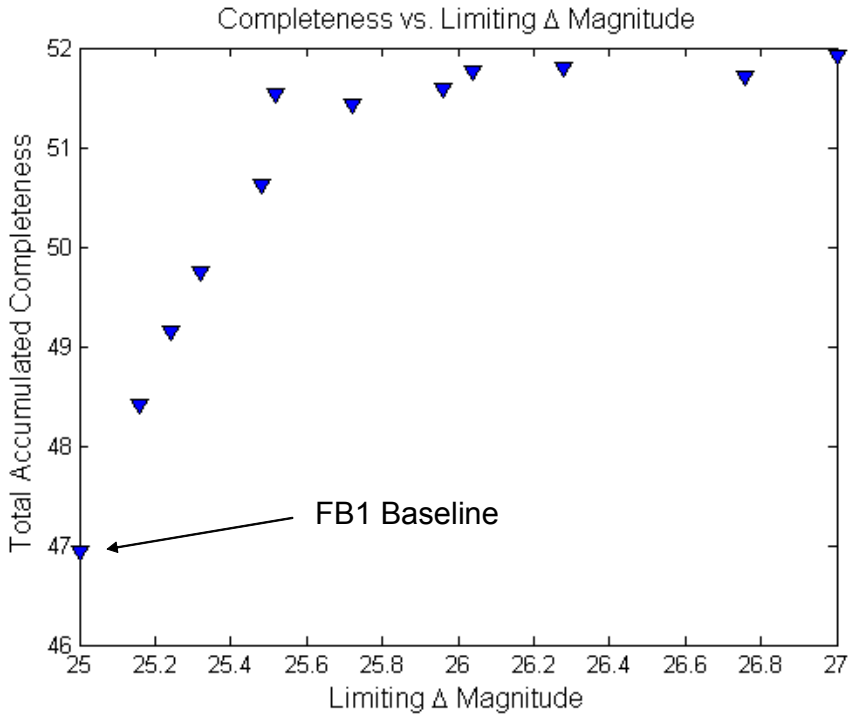

Fig. 16. Completeness vs. limiting Δmag. The baseline is set at Δmag = 25.

Cumulative completeness increases rapidly with an increase in limiting Δmag and then levels off at Δmag ~ 26. Observing to deeper limiting Δmag requires more time, pushing the maximum point of observation further to the right on the stellar completeness curves (Fig. 7). As the limiting Δmag is increased, the curves begin to level off (like the low luminosity completeness curve in Fig. 7a). At higher Δmag, more curves begin to exhibit the same shape of the low luminosity curves and there is little benefit gained from observing deeper (the integration time is cut from the flat part of the curve). If the baseline limiting Δmag requirement increases to Δmag = 25.7, the cumulative completeness increases by ~9%. It is possible with the current error budget to increase to Δmag = 25.7 (Shaklan et. al., 2005) but this would involve instituting much more stringent stability requirements to achieve the same margin.

Exo-zodi brightness, throughput, IWA and limiting Δmag all have a strong effect on system requirements. Overhead, on the other hand, is not a strong driver. There are 2 types of instrument overhead, fixed and proportional. Fixed overhead includes dither and roll maneuvers, re-pointing and settling time (independent of star brightness) while proportional overhead includes wave front sensing and control. Here we consider only fixed overhead. Figure 17 is a plot of the number of HZs searched vs. overhead.

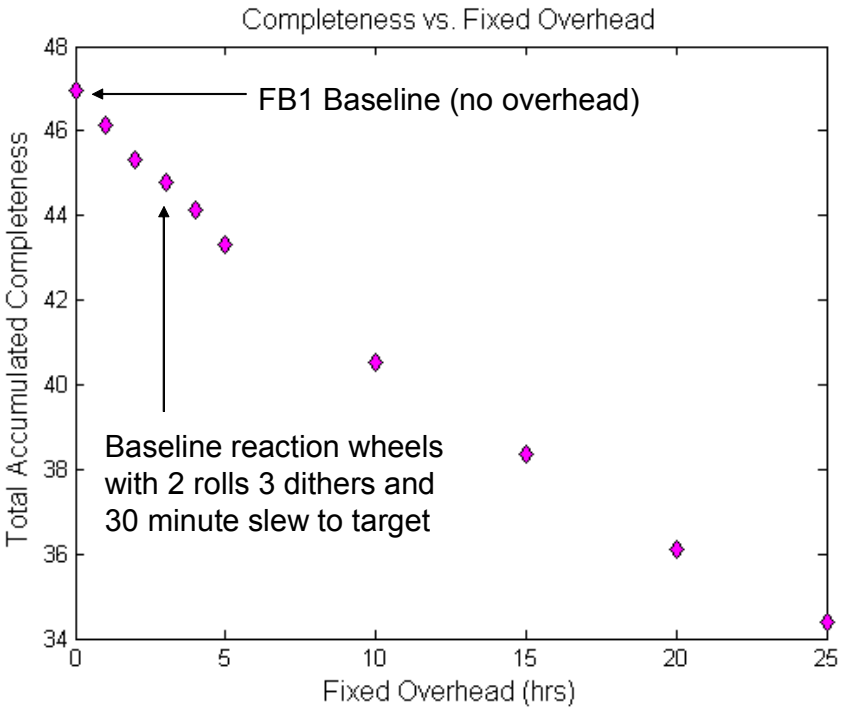

Fig. 17. Comparison of completeness at different overhead times. The baseline is set at zero hours of overhead. The current baseline reaction wheels give 2 hours of overhead.

Overhead is driven by choice of the reaction wheels and instrument mechanical damping. With smaller reaction wheels, the telescope is not able to change orientation as rapidly and overhead increases. As one may expect, this decreases the cumulative completeness as shown in Fig. 17. The maximum number of visits to a single star also decreases with an increase in overhead. For long overhead times (15-20 hours), the maximum number of revisits is 6 (as opposed to 24 for zero hours of overhead). However, since the completeness does not drop off rapidly, the loss due to overhead is not a strong driver on the system requirements. Completeness variations due to the changes in overhead and the aforementioned system parameters are listed in Table 2.

	Exo-Zodi Brightness	Limiting ΔMag	Throughput	IWA	Overhead	Number of Targets	Total Accumulated Completeness	Percent Change
	x Earth		%	mas	hrs			
FB1 Baseline	1.5	25	10.8	65.5	0	88	46.90	0%
Case #1	1	25	10.8	65.5	0	92	49.06	4%
Case #2	4.5	25	10.8	65.5	0	72	38.60	-21%
Case #3	5.0	25	10.8	65.5	0	69	37.50	-25%
Case #4	1.5	25.48	10.8	65.5	0	91	50.63	7%
Case #5	1.5	25.72	10.8	65.5	0	92	51.42	9%
Case #6	1.5	25.96	10.8	65.5	0	91	51.59	9%
Case #7	1.5	25	5.4	65.5	0	66	36.58	-28%
Case #8	1.5	25	25.0	65.5	0	113	60.17	22%
Case #9	1.5	25	30.0	65.5	0	118	62.94	25%
Case #10	1.5	25	10.8	55	0	105	57.82	19%
Case #11	1.5	25	10.8	42.5	0	126	70.57	34%
Case #12	1.5	25	10.8	28.4	0	137	79.53	41%
Case #13	1.5	25	10.8	65.5	2	84	45.31	-4%
Case #14	1.5	25	10.8	65.5	10	79	40.52	-16%
Case #15	1.5	25	10.8	65.5	15	76	38.35	-22%

Table 2. Case studies for changing system requirements and exo-zodi level. The baseline is shown in blue and the pink boxes indicate changes to the baseline. The boxes in the last column represent positive (green) and negative (red) relative changes to the baseline.

5. CONCLUSIONS

Program completeness has been optimized with one technique based on the single visit efficiency threshold optimizations. The quintessential result of this optimization is that some stars are visited many more times than previously anticipated. This is an unexpected result that arises because the marginal gain due to orbital motion of the observable habispheres provides higher completeness even after 2-3 years than would be gained from visiting faint, distant stars. With the baseline assumptions, on average there are approximately 7 visits mandated for each visited star. Interestingly, there is significant variation in the number of visits per star as some stars are only observed once while others are observed upwards of 20 times. This type of behavior is not significantly affected by changing the visit spacing or minimum revisit time, which justifies the freezing and discretization assumptions for this iteration of program completeness.

The effect of system parameters on cumulative completeness was also explored. Overhead was found to not be a strong driver (greater than 20% increase or decrease in cumulative completeness) for changing the system requirements (an increase of over 12 hours or a factor of 6 is required to have a significant effect). Increasing the limiting Δmag (to Δmag=25.7 from Δmag=25) improved detections by 9%. In addition, increasing exo-zodi brightness is a large driver in decreasing the science benefit at a change of a factor of 5. In contrast, increasing throughput (to 21% from 10.8%) or decreasing the inner working angle (to 55 mas from 65.5 mas) would create a strong science benefit and thus are strong drivers of the system architecture.

References:

Brown, R. A. "Buy-Back Auction optimization of TPF-C exposure times for first-visit completeness," (memo) May 23, 2005.

Brown, R. A. "Obscurational Completeness," ApJ **607**:1003-1013, 2004.

Brown, R. A. "Photometric and Obscurational Completeness," ApJ, 2005.

Brown, R.A. "Whole Mission Simulation – First Cut_2," (memo) July 9, 2005.

Hunyadi et. al. "Single Visit Completeness Auction Memo," October, 2005.

Shaklan, S.B., et. al. "The Terrestrial Planet Finder Coronagraph Dynamics Error Budget," Proc. SPIE Vol. 5905 (2005).

Shaklan, S. B. and Lisman, P. D. "TPF Coronagraph Design Reference Mission Study," Proc. TPF SWIG February 25, 2004.

STDT. "Terrestrial Planet Finder Coronagraph Science Requirements," (memo) September 8, 2005.