

# Searching for transits in data with long time baselines and poor sampling

B. Tingley<sup>1</sup>

Instituto de Astrofísica de Canarias, C/ Vía Láctea, s/n, E38205 – La Laguna (Tenerife), España  
e-mail: btingley@iac.es

Received 7 October 2010 / Accepted 7 February 2011

## ABSTRACT

**Aims.** The standard method of searching parameter space for transits is ill-suited to data sets with long time baselines and poor temporal coverage, such as that anticipated from Gaia. In this paper, we present an alternative method for identifying transit candidates in such data, one focusing on finding periodicity in high S/N outliers.

**Methods.** We describe a technique for testing a small number of flux measurements for periodicity and consistency with an origin in a transit with a constant change in flux and test their performance with Monte Carlo simulations. To complement this, we also include a description of a statistical method to analyze the distribution of these measurements to determine if they are normally distributed around a constant, reduce flux consistent with a planetary transits.

**Results.** Large numbers of light curves can be quickly scanned for transit signatures with minimal loss in effectiveness for data sets with long time baselines and poor temporal coverage, where one observation per transit is the norm by testing for periodicity and analyzing their distribution.

**Conclusions.** If the noise characteristics of the data set and the intrinsic noise of the individual stars are understood, this method focusing on statistical outliers is nearly equivalent to the standard method of scanning parameter space and significantly faster, if the signal  $\gg$  noise, the individual transits are sampled no more than once and a periodicity test is employed. Moreover, the test for a transit origin can eliminate additional false positives.

**Key words.** methods: statistical – planetary systems

## 1. Introduction

Gaia is a space mission that will be launched in November 2012. A cornerstone mission of the ESA Space Program, it is designed to be the successor to the highly successful *Hipparcos* mission. It is primarily designed to perform astrometry with unprecedented precision on up to 1 billion stars in our Galaxy, gathering photometry and spectroscopy simultaneously to enable a detailed picture of Galactic kinematics. As part of accomplishing these tasks, Gaia will obtain on the order of 100 photometric measurements for each of these stars during its 5-years mission. These measurements will have a precision ranging from a few  $10^{-3}$  mag for stars with  $V < 17$  to a few percent down to  $V \sim 20$  with a semi-irregular temporal sampling comprised of loose clusters of a few observations over 1 or 2 days (although never closer than  $\sim 0.08$  d) widely spaced in time (Jordi et al. 2006). Obviously, such a vast data set presents myriad opportunities for secondary science. The focus of this paper will be to describe a technique for identifying planetary transit candidates in Gaia photometric time series.

Considering the characteristics of Gaia photometry, we can conclude that hot Jupiter-red dwarf systems offer the best prospects for planet discovery. A transit of a hot Jupiter around a solar-like star will result in a decrease in brightness of approximately 1%. As each transit is observed singly, it must be observed with a high S/N to be detected. Given the limits of the photometry, individual transits of even the brightest solar-like stars would be detected at only the  $2-3\sigma$  level. Red dwarfs, however, are physically smaller than solar-like stars. As transit

depth goes approximately as the radius of the planet divided by the radius of the star squared, red dwarf transits would therefore be correspondingly deeper, ranging from a few percent deep for the largest red dwarfs to 10% or more for the smallest. This allows transits to be detected around stars for which the time series have much lower precision, corresponding to fainter stars in the Gaia data set, which are far more numerous. Moreover, red dwarfs are the most common stars in the Galaxy in a distance (rather than magnitude) limited sample – in fact, there may be up to  $10^8$  red dwarfs in the portion of the Gaia photometric data set of interest, i.e. those stars with  $V < 17$ . These stars would have easily detectable ( $>3\sigma$ ) individual transits and, being relatively nearby, precise distances and absolute magnitudes from the Gaia primary science mission, allowing the immediate elimination of many false positives.

Red dwarfs are otherwise highly underrepresented targets in transit searches and indeed in exoplanet searches in general due to their redness (which limits precision of radial velocity measurements using spectra, as most spectrographs operate in the visible where red dwarfs have relatively few lines), their faintness, and their low sky density at magnitudes open to radial velocity searches. A few projects have been begun that focus on discovering planets around red dwarfs (RoPACS, surveying  $\sim 10^4$  red dwarfs in the infrared in several fields totaling several square degrees<sup>1</sup> and the MEarth project, monitoring the brightest  $\sim 2000$  red dwarfs individually with an array of robotic telescopes (Irwin et al. 2009), but these target a relatively small

<sup>1</sup> <http://star.herts.ac.uk/RoPACS/index.html>

sample of stars in an attempt to discover rocky planets. Unless such rocky planets orbiting red dwarfs are far more common than giant planets around solar-like stars, these projects will have to be fortunate to make more than a single detection. Clearly, the case for Gaia will be different – while rocky planets are likely beyond its capabilities, the potential harvest of hot Jupiters is enormous, unless they are extraordinarily rare around red dwarfs – contingent upon our ability to detect them in the light curves and eliminate a large fraction of false positives.

Our approach focuses on two characteristics that are unique to planetary transits: periodicity and an essentially constant drop in flux during transits, which we will refer to as flatness. By selecting light curves that have an excess of outliers corresponding to a reduction in flux and then testing these outliers for periodicity and flatness, it should be possible to eliminate abnormally noisy stars (whether from bad data or activity) and identify those stars that exhibit events consistent with grazing eclipsing binaries and exoplanets.

## 2. Method

Two different styles of transit searches and associated algorithms are currently discussed in the literature – dedicated transit surveys for which single fields are sampled nearly continuously, ultimately collecting thousands of observations (e.g. Jenkins et al. 1996), and serendipitous approaches that comb photometric surveys gathered for other purposes which contain  $\sim 10$  observations and endeavor to identify candidates based on a single observation during an eclipse/transit (e.g. Ford et al. 2008). Gaia, which will obtain (usually) only one observation per transit but will detect multiple transits in many cases, falls somewhere in between. It therefore follows that an optimal transit detection algorithm for Gaia photometry should employ aspects of both. With this as a starting point, a brief overview of the two techniques is instructive.

The goal of the continuous approach is to observe individual transits many times in an effort to build up statistical significance, leading to confident detections of transits that have a signal on the same order as the noise. To perform this task, a large set of “test” light curves covering parameter space is generated and compared to the observed light curve. A numerical method – such as the BLS (Kovács et al. 2002) – is used to assess how well each individual test light curve matches the observed light curve, with the best match designated the most statistically likely set of transit parameters. The number of possible test transits is reduced by using the properties of the transits (a periodic event that can be well-approximated by a constant decrease in flux for a fixed duration) and by taking small enough steps sizes (particularly in period) that the correlation coefficient between adjacent test light curves is 0.5 (Jenkins et al. 1996) – or some similar criterion.


Positives and negatives exist in a straightforward application of this approach to Gaia light curves. On the positive side, this approach is very effective in eliminating random or non-periodic events due to the implicit periodicity. Additionally, it will assumedly recover 100% of transits in the data, although this comes at a high price computationally. That is the main downside of this approach: the number of test light curves necessary to cover parameter space increases linearly with the time baseline (Jenkins et al. 1996). Given Gaia’s long time baseline and considerable number of targets, this approach becomes computationally ponderous in the extreme. Moreover, Gaia will in the vast majority of cases observe each transit only once – a fundamental difference between Gaia photometric time series and the

principles upon which this approach is constructed. Ideally, we would like to derive an approach that preserves the sensitivity to periodicity – and the ability to filter out false positives that it brings – while being significantly lighter computationally.

The serendipitous approach is completely different. Designed to identify single outliers in millions of light curves containing a dozen or so observations, the main challenge these algorithms face is separating out true “signal” outliers from statistical/noise outliers. Dupuy & Liu (2009) compare and contrast several different detection algorithms, stating that the goal is to identify the algorithm that is best at detecting faint fluctuations while simultaneously rejecting false positives – generally considered to be dominated by variable/active stars. All of these algorithms have the same basic principle, however: identify stars with a small number (usually one) of anomalous outliers and label them candidates, particularly if the other fluxes have a small variance compared to the depth of the outlier. The advantage to this approach is that it is very light computationally and can thus be readily used to analyze large numbers of stars. Its weaknesses are that it is vulnerable to both astrophysical and serendipitous (e.g. a dead pixel or bad column) false positives and that it provides very little information about the event besides some information about the activity level of the star. This fact is not crucial in the serendipitous approach, as they anticipate that only a single transit will be observed once and therefore it is impossible to do more. We, however, expect to discover so many candidates in the Gaia photometric data set with multiple in-transit observations that we can afford to ignore those with just a single outlier. The goal here is different: we have to go beyond mere detection and focus instead on identifying the best candidates – this requires extracting additional information from the light curves to establish candidacy.

An examination of these two techniques does however suggest how to proceed. It is far lighter computationally to look for high S/N outliers corresponding to a decrease in brightness rather than to perform a series of  $\chi^2$  tests on millions of possible test light curves. This is quickly done by simply sorting the fluxes in the light curve, which should gather all of the most significant outliers together. Then, these observations need to be tested for the primary two qualities of transits: periodicity and flatness – plus an additional test for activity, such as the one described by Dupuy & Liu (2009). In the process of doing this, we additionally have to determine where the outliers end and the extreme values from the continuum begin.

Stellar activity is another issue that must be addressed. It is possible for active stars to mimic all manner of signals when sparsely sampled, including those produced by transiting exoplanets. Our approach will be to screen the candidates by calculating the depth of the presumed transit in units of signal-to-noise by calculating the  $\sigma_{\text{rms}}$  of the apparent out-of-transit flux measurements to the mean of the presumed in-transit flux measurements, in addition to magnitude or flux units. Most active stars are varying constantly, so they would typically exhibit a relatively high out-of-transit  $\sigma_{\text{rms}}$ , reducing the statistical significance of any transit-like event detected. We anticipate that requiring both a high signal-to-noise and a transit depth consistent with a transiting exoplanet for candidacy will allow us to eliminate the vast majority of false positives arising from variable or active stars. This would leave only transiting exoplanets and two classes of eclipsing binaries: those with grazing eclipsing or those occurring in high mass ratio systems. We hope that many of these can be eliminated based on Gaia radial velocity measurements, astrometric variations, and absolute magnitudes.



**Fig. 1.** Sample light curve used in description. It has a period of 1.1 days and a duration of 0.1 days, with a transit depth of 0.025 mag and white Gaussian noise with a  $\sigma_{\text{rms}}$  of 0.005 mag. The event is therefore  $5\sigma$  deep.

### 2.1. Periodicity test

The periodicity test is most crucial one to this approach to transit detection and also the most difficult to construct. Part of the reason for this difficulty is that a series of individual observations of different transits of single star is not completely periodic, despite the fact that the transit itself is. The reason for this quasi-periodicity is that a single observation of a transit can occur any-time during the transit and we don't have any a priori knowledge of when. This is a very significant distinction, requiring a very different approach than simple, strict periodicity. The approach and implementation are described here.

The approach is constructed around the concept that the temporal difference between any two in-transit observations ( $\Delta_{ij}$ , where  $i \neq j$ ) is an integer multiple ( $m$ , the transit number) of the period ( $\tau$ ), plus or minus the duration ( $D$ ) of the transit, assuming box-shaped transits – it is trivial to convince oneself of this – with the logical extension that  $m = 0$  if both observations are in the same transit. Extending this, we claim that the maximum and minimum possible periods are for a given  $m$ ,  $\Delta_{ij}$  and  $D$  are:


$$\tau_{\text{max,min}} = \frac{1}{m}(\Delta_{ij} \pm D). \quad (1)$$

From this equation, we can infer that the exact period cannot be found, only a period restricted to within  $\pm D/m$ , which can be quite small for large  $m$ . Moreover, we can limit  $m$  by setting a minimum possible period (for example, 0.75 days, slightly less than the shortest known exoplanetary period), reducing the number of possible periods – in theory,  $m$  could be infinite, but that would correspond to an aphysical  $\tau = 0$ . With this, we can derive an expression for  $m_{\text{max}}$ :


$$m_{\text{max}} = \frac{\Delta_{ij} \pm D}{\tau_{\text{min}}}. \quad (2)$$

As a consequence of this, we can see that a large  $\Delta_{ij}$  will have a large  $m_{\text{max}}$ , meaning a large number of possible periods but with a small allowed range ( $\pm D/m$ ), while a small  $\Delta_{ij}$  will have a smaller  $m_{\text{max}}$ , meaning fewer, but less constrained, possible periods.

While this is no solution, it does offer insight into how a solution may be found: given a set of  $N$  observations to test for transit-like near-periodicity (see Fig. 1), the temporal differences between all of the observations can be calculated and all possible



**Fig. 2.** Example of temporal difference between observations. This figure shows the temporal differences between a set of 5 in-transit observations with the associated allowed period ranges based on the test duration  $D$  and the transit number  $m$  depicted as error bars. The true period is 1.1 days, marked by the dotted line, and the test duration is 0.2 days. The size of the temporal differences ( $\Delta_n$ ) increases from left to right. Notice how the period ranges narrow during this progression, due to the increasing transit number. By testing the periods defined by the smallest temporal difference ( $\Delta_0$ ), which has the smallest  $n$  and therefore fewer possible periods, then the same for the next-smallest temporal difference ( $\Delta_1$ ), etc., we can most efficiently converge on the best period. The  $\Delta_n$ s in this case are 4.4979, 7.4797, 80.9989, 88.4786, 117.0229, 121.5207, 124.50251, 129.0004, 205.5014, and 209.9993. The corresponding values of  $m$  that yield the correct period are 4, 7, 75, 81, 107, 111, 114, 118, 188, and 192.



**Fig. 3.** Temporal difference between observations, zoom. This figure is the same as above, but focused on the solution, with the crosshatched region showing the constraints on the period above and beyond the limit bars. Here, one can see an example of a branch in the possibilities at  $\Delta_1$ , as two possible periods from  $\Delta_2$  fall within the period range described by  $\Delta_1$ . This is a relatively common occurrence, so the code has been designed to follow all possibilities should a branching occur.

periods with limits determined (see Fig. 2), where the  $\Delta_{ij}$ s are depicted from smallest on the left to largest on the right and the error bars represent the possible period rangers for each  $\Delta_{ij}/m$  (as defined by Eq. (2)). It only remains to find a single period that can satisfy all of the  $\Delta_{ij}$ s. For the purpose of simplicity, we sort the  $\Delta_{ij}$  from smallest to largest and recast them as  $\Delta_n$ s, such that the smallest  $\Delta_{ij}$  would be  $\Delta_0$ , the next smallest  $\Delta_{ij}$  and so on, with the largest  $\Delta_{ij}$  defined to be  $\Delta_M$ .

The simplest approach would be to use brute force to find the period, testing each possible period from the  $\Delta_M$  to see if

a matching possible period exists for all of the other  $\Delta_n$ s. This is, however, tediously slow, even for a relatively small data set. It is far more efficient to start with  $\Delta_0$ , which would have the lowest  $m_{\max}$  – i.e. the fewest possible  $m$ . For each  $m$ , the period range is established:  $\tau_m = \Delta_0/m \pm D/m$ . Then  $\Delta_1$  is examined to see if any of its period ranges overlap. If one or more do, the allowed period ranges are restricted again, with the maximum allowed period being the smallest of  $\Delta_0/m_0 + D/m_0$  and  $\Delta_1/m_1 + D/m_1$  and the minimum allowed period the largest of  $\Delta_0/m_0 - D/m_0$  and  $\Delta_1/m_1 - D/m_1$ . This same process is then applied to  $\Delta_3$ , again restricting the periods, then continuing through  $\Delta_4$  and, should they all continue to successfully overlap,  $\Delta_M$ , where  $M = N(N - 1)/2$ . Should one of the  $\Delta_n$ s fail to have a period that falls into the period range, then that particular  $m_0$  does not correspond to the period. Should none of the  $m_0$ s produce a period range for which all of the other  $\Delta_n$ s have period ranges that overlap, then the in-transit observations are non-periodic. Only one significant wrinkle requires special attention: it is possible that multiple values for  $\Delta_{k+1}/m_{k+1}$ s will nest inside the period range allowed by a single value of  $\Delta_k/m_k$ . The result of this is a branch in the possibility tree – and all of these branches must be followed for the proper period to be found in all periodic cases.

In practice, this is not entirely sufficient to solve the problem, due to the interplay between transit duration and period as large durations combined with short periods will tend to dominate, even though they correspond to situations that one would never encounter in nature. It is therefore necessary to restrict the maximum test duration to reflect what is known about the relationship between planetary transit durations and periods, as restricting the highest test duration reduces the false positive probability. From Tingley & Sackett (2005), we know that  $D = \alpha\tau^{1/3}$ , all other factors (such as stellar mass, stellar radius, and planetary radius) being equal, with  $\alpha$  defined as the constant of proportionality. The highest such  $\alpha$  for all known exoplanets is about 3.1 (with  $\tau$  in days and  $D$  in hours) for HAT-P-7b (Pál et al. 2008), while the average is about 2. The best way to include this effect into the code is to test different durations, each of which will have its own  $\tau_{\min}$ . These  $\tau_{\min}$  should follow the relation mentioned above; to be rigorous, one can also use different  $\alpha$ s, if one desires. We suggest using a mean value ( $\alpha = 2$ ) for the maximum test duration – or perhaps an even smaller when searching for transits in red dwarfs – as, according to Tingley & Sackett equation,  $\alpha$  should go approximately as  $(R_{\text{planet}} + R_{\star})M_{\star}^{-1/3}$ . The extreme case HAT-P-7b is at least three times larger than a typical red dwarf with  $R_{\star} \sim 1.8 R_{\odot}$ . While its mass is similarly large ( $M_{\star} \sim 1.5 M_{\odot}$ ), the dependence of  $\alpha$  on stellar mass is much weaker. Moreover, the planet itself is exceptionally large ( $R_{\text{planet}} \sim R_J$ , inflated by its proximity to such a large, comparatively hot ( $T_{\text{eff}} = 6350$  K) star, which also increases  $\alpha$ ). By contrast, the two known transiting red dwarf planets GJ436b (Bean et al. 2008) and GJ1214b (Charbonneau et al. 2009), both of which are significantly smaller than Jupiter, have  $\alpha$ s of  $\sim 0.4$  and  $\sim 0.55$  respectively.

It is necessary to allow for the possibility that, according to the Gaia scanning law, a single transit might be observed multiple times. It is a relatively simple matter to adapt the approach described above for this purpose: just throw out any  $\Delta_n$  less than the minimum allowed period. While some small fraction of the information is lost, the code will run normally otherwise and affect on recoverability is very small.

Another capability that can be built into the periodicity test is the ability to exclude a true statistical outlier that may interpose itself into the set of in-transit observations. This one “interloper”

(as we will refer to them) has a good chance of causing the period test to fail. This will not be an uncommon occurrence, as one out of every  $\sim 740$  observations will be a  $3\sigma$  outliers on the faint side of the distribution assuming that the noise is Gaussian – in other words, one out of five Gaia light curves. So unless one wishes to dismiss erroneously a sizable fraction of the total candidates in the data set, the code must be capable of excluding a single interloper and finding the proper period despite it.

Monte Carlo simulations are perhaps the most rigorous method for testing this technique. By creating large numbers of simulated light curves with different periods and durations and random observation ephemerides, it is possible to test the response of this algorithm to realistic conditions. By tweaking the code, it is easily possible to force a specific number of randomly chosen in-transit observations (using the transit period-duration relation described earlier), including zero in-transit observations (completely random), which provides a necessary baseline for comparison purposes.

## 2.2. Flatness test

Given the nature of transits, the in-transit observations should be consistent with an event of more or less constant depth. If this is not the case, it is likely that the detected event arises from some cause other than a transit, for example a poorly sampled eclipsing binary, an active star, or simply unexpected systematic noise. It is possible to derive a test that can be quickly and easily performed, again making the assumption that the noise is Gaussian distributed. If the observations are truly caused by a transit of constant depth, the fluxes should be Gaussian distributed around the mean – the measured depth of the transit; however, if these observations are a collection of statistical outliers, the distribution should be different, comprising the tail end of the statistic distribution centered around the continuum value. Therefore, a determination of the sample variance of the in-transit observations should be an effective test of the “flatness” of the in-transit subset:

$$\sigma^2 = \frac{1}{N-1} \sum_1^N (x_i - \bar{x})^2 \quad (3)$$


where the  $x_i$  are the members of the subset of the  $N$  faintest observations, assuming the precision per observation is approximately constant.

Using Monte Carlo simulations, it is possible to evaluate the effectiveness of this test, calculating the mean (the event depth) and the sample variance of the  $N$  faintest points drawn out of a sample of  $T$  randomly generated, Gaussian distributed points for a large number of trials (in this case,  $10^7$ ). This can then be compared to the sample variance of  $T$  Gaussian distributed points representing the transit, with the understanding that the transit will have its own constant event depth that separates it from the continuum.


## 3. Results

We performed Monte Carlo simulations for the periodicity test, allowing the number of in-transit observations to range from 5 to 17, varying the number of non-periodic interlopers from 0 to 2, using two different values for  $\alpha$  (2 and 3) and with the capability to exclude non-periodic interlopers both enabled and disabled. The results of these simulations are shown in Fig. 4. In general, these results show that the periodicity test finds periodicity in

B. Tingley: Searching for transits in data with long time baselines and poor sampling



**Fig. 4.** Results of Monte Carlo simulations of Periodicity test. These figures show the results of the Monte Carlo analysis of the periodicity test, yielding the probability that a sample of in-transit observations will pass under a wide range of conditions. We disable (thick lines) and enable (thin lines) the ability to exclude a single interloper for a periodic (continuous lines) and non-periodic (dashed lines) set of in-transit observations. The size of this set is allowed to range from 5 to 17, with the number of interlopers ranging from 0 (*top row*) to 2 (*bottom row*) for two different  $\alpha$ s, 2 (*left column*) and 3 (*right column*). Note first and foremost that non-periodic samples exhibit a significant probability of passing the periodicity test if the number of observations is small. This suggests a lower limit to the number of in-transit observations to be confident in the periodicity. This limit has an additional dependence on  $\alpha$  as well. Notice that enabling interloper exclusion increases this lower limit by approximately two – effectively one, as the interloper is removed from consideration. In addition, if interloper exclusion is not enabled and an interloper is present, the probability that the periodicity will still be recovered, regardless of the number of true periodic observations, drops significantly: about 18% for  $\alpha = 3$  and only 10% for  $\alpha = 2$ . Lastly, if two interlopers are present among the in-transit observations, even with single interloper exclusion on the probability that the sample will pass the test is low: 34% for  $\alpha = 3$  and 20% for  $\alpha = 2$ . Such situation should be rare, however, unless the noise distribution is very non-Gaussian.



**Fig. 5.** The variance of the variance of the in-transit observations: Gaussian outliers vs. true transits. These figures show how the variance of the presumed in-transit observations varies depending on total number of observations ( $T$ ), number of in-transit observations ( $N$ ) and event depth in S/N ratio for true transits or a collection of Gaussian outliers. The dark gray region defines the  $1\sigma$  range of the possible sample variances of true in-transit observations, while the light gray region define a similar range for presumed in-transit observations that are actually Gaussian outliers, with the solid lines marking the center of these distributions. The dotted line marks the probability that a “false positive” transit of the specified depth by  $N$  Gaussian outliers out of a sample of  $T$  observations could be randomly generated. As an example, let us assume that a transit with a depth corresponding to a signal-to-noise ratio of 2.5, comprised of 12 in-transit observations out of 200 total observations. The  $\sigma_{\text{rms}}$  of the out-of-transit observations and the in-transit observations is then calculated and compared ( $\sigma_{\text{in}}/\sigma_{\text{out}}$ ). Looking at the appropriate plot, we can see that a value higher than 0.5 or so is highly unlikely to be caused by chance collection of statistical outliers based simply on the distribution of the fluxes of the observations, without regard to periodicity.

essentially all of the periodic samples. If a non-periodic interloper is present, the probability that the test will find periodicity drops by approximately an order of magnitude, unless the option to remove interlopers is enabled, in which case the probability returns to 100%. Enabling interloper exclusion comes with a price, however; the probability that a non-periodic sample will pass the periodicity test increases – a false positive. Without interloper exclusion enabled, the number of in-transit observations must be at least 7 or 8 (depending on  $\alpha$ ) before the false positive probability drops below 1%. If interloper exclusion is enabled, then two more in-transit observations are required to drop the false positive probability below this level – although one of these is the interloper, so the actual requirement to maintain confidence level is one additional in-transit observation.

The results of the Monte Carlo simulations of the flatness test are shown in Fig. 5. In these figures, the variance of the presumed in-transit observations is shown to be fairly effective at differentiating between the events generated by true transits compared to those that are randomly generated via Gaussian

noise. They also emphasize how rare it is for a randomly-generated event to produce an even moderately deep transit. The test does become less effective as event depth increases; however, the chance that such an event could randomly occur also becomes much lower. The test becomes more effective as the number of in-transit observations increases, as this acts to reduce the variance for a given event depth for the randomly generated events and the variance for both. It becomes less effective as  $T$  increases, as this increases the depth of extreme randomly-generated events. But again, increasing  $T$  generally increases the number of in-transit observations, which follow Poisson statistics, given a probability of being in transit ( $D/\tau$ ) and  $T$ .

These Monte Carlo simulations by themselves are too cumbersome to be run for each and every candidate. Instead, it is more practical to tabulate confidence values based on the average sample variance and the associated variances in advance and interpolate to estimate confidence levels. However, for individual interesting candidates, full simulations (perhaps with a better

noise model, which can be determined through an analysis of the data a posteriori) can be performed easily and relatively quickly.

#### 4. Hypothetical project: transit survey of red dwarfs

The apparent utility of this approach to transit detection encourages the application of a similar strategy from ground. Surveys for transits around red dwarfs have been frustrated by the scarcity of such stars above magnitude 17 in a typical field of view (see Fig. 6). Therefore, large portions of the sky need to be surveyed to accumulate the number of red dwarfs needed to ensure a reasonable chance of detecting at least one transiting planet, unlike the earlier main sequence stars targeted by normal surveys. As this technique requires high S/N events, the small planets that are typically the goal of red dwarf transit surveys (e.g. Nutzman & Charbonneau 2008) would still be mostly out of reach. It would however be possible to establish the existence of hot Jupiters in such systems – a logical first step.

During this hypothetical survey, a field would be observed for 15 min three times a night, each of these observing blocks separated by 2.5 h, which is comfortably longer than a hot Jupiter in a circular 5-day orbit around an M0 star – most hot Jupiters have circular or nearly circular orbits. This ensures that any transit captured will occur during a single 15 min block, which is what the technique described in this paper is designed for. 10 fields could be sampled 3 times each night in this fashion. Figure 7 shows the resulting window function that arises from three months of telescope time, using a cut-off of 5 in-transit observations as sufficient to recover the period with greater than 95% confidence, assuming that the 10 fields are observed in a different order each night to reduce aliasing. Over half of the giant planets transiting red dwarfs with a period of 4 days and virtual all with periods less than 2 days would be detected during such an observing program, despite only 270 samples per star – an order of magnitude fewer than a typical transit survey. The newest generation of survey telescopes will be capable of monitoring large numbers of red dwarfs in this manner; for example, with Skymapper (1.35-m aperture, 5.7 square degree field), some 60 000 red dwarfs would have good enough photometry for giant planet detection. This is a robust number, as transiting hot Jupiters are detected at a rate of about 1 per 10 000 main sequence stars surveyed (Bayliss et al. 2009). If this rate holds, we would expect to discover several hot Jupiter + red dwarf systems in such a survey.

#### 5. Conclusions

In the case of poorly sampled data with noise that can be described as truly Gaussian or nearly Gaussian, the approach to the detection of singly-sampled, high signal-to-noise transits described in this paper functions adequately, capable of recovering essentially 100% of the test light curves from the periodic sample. It does, however, discover periodicities for a significant percentage of non-periodic (i.e. random) sources; it is not clear, however, that any other technique could avoid this problem, including the standard approach of folding the light curve at different test periods and testing different transit ephemerides and durations, as described in e.g. Bordé et al. (2003). Additionally, this approach is significantly lighter computationally than the standard approach, performing the necessary tests in several orders of magnitude less time. The code

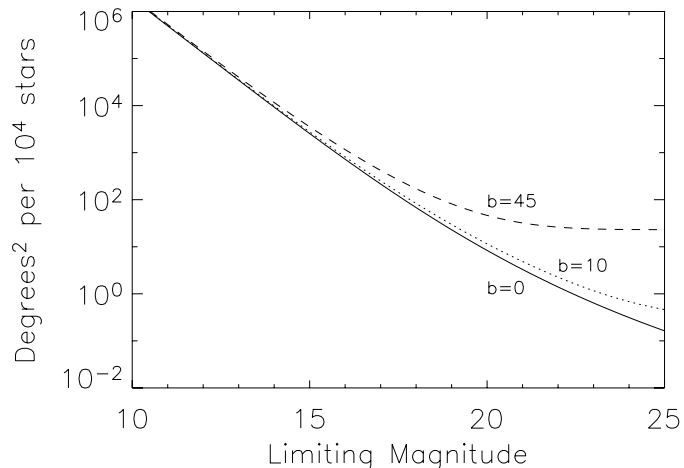


Fig. 6. Number of M dwarfs per square degree as a function of V magnitude and Galactic latitude (Arentoft, priv. comm.).

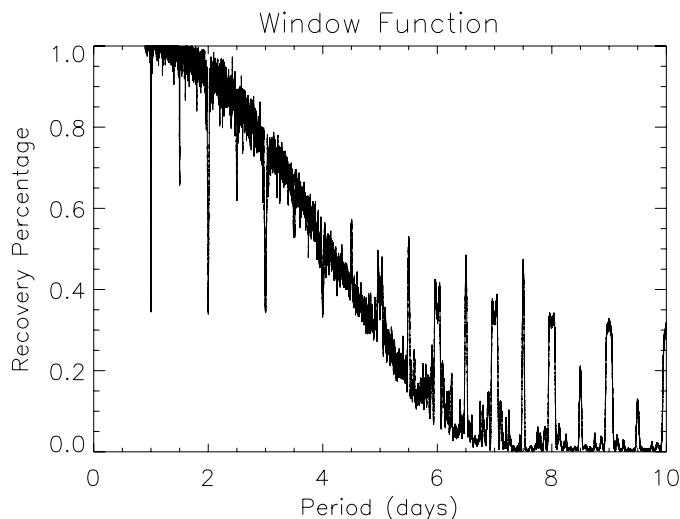


Fig. 7. Window function for a hypothetical red dwarf transit survey. This figure depicts a typical window function describing the probability of detection for hot Jupiters transiting red dwarfs as a function of period for a survey where the fields are sampled for 15 minutes three times a night for three months.

can be slightly modified to deal with the presence of a single, non-periodic interlopers in the periodic in-transit sample. This comes at a cost, however, as the rate of false positives increases.

This approach to transit detection codes suggests a new style of ground-based transit survey, one focused on relatively deep planetary transits in field red dwarfs. Instead of acquiring several thousand observations of one or two fields in order to detect low-amplitude events, this analysis opens the possibility of surveying many fields, gathering 200–300 observations for each target, enabling confident detection of events 0.05 to 0.1 mag deep. Instead of surveying at or near the Galactic plane, the Galactic polar region would be the best part of the sky to survey, as red dwarfs are nearly isotropically distributed down to limiting magnitude deep enough for the purposes of the survey – in general, they are within a kiloparsec down to  $V = 18$ , which corresponds well to the scale height of the Galactic disk (500 pc). This means

that the number of field red dwarfs would be approximately the same in all directions, while the number of background eclipsing binary stars capable of blending with a target star to produce an event that mimics a planetary transit would be vastly reduced, simplifying the necessary follow-up observations for confirmation.

## References

- Bayliss, D. D. R., Wel Drake, D. T. F., Sackett, P. D., Tingley, B., & Lewis, K. 2009, *AJ*, 137, 4368
- Bean, J. L., Benedict, G. F., Charbonneau, D., et al. 2008, *A&A*, 486, 1039
- Bordé, P., Rouan, D., & Léger, A. 2003, *A&A*, 405, 1137
- Charbonneau, D., Berta, Z. K., Irwin, J., et al. 2009, *Nature*, 462, 891
- Dupuy, T. J., & Liu, M. C. 2009, *ApJ*, 704, 1519
- Ford, H. C., Bhatti, W., Hebb, L., et al. 2008, in *Classification and Discovery in Large Astronomical Surveys*, ed. C. A. L. Bailer-Jones (Melville, NY: AIP), AIP Conf. Proc., 1082, 275
- Hébrard, G., Robichon, N., Pont, F., et al. 2006, in *Tenth Anniversary of 51 Peg-b: Status of and prospects for hot Jupiter studies*, ed. L. Arnold, F. Bouchy, & C. Moutou, 193
- Irwin, J., Charbonneau, D., Nutzman, P., & Falco, E. 2009, in *Transiting Planets*, ed. F. Pont, D. Sasselov, & M. Holman (Cambridge: Cambridge Univ. Press), IAU Symp., 253, 37
- Jenkins, J. M., Doyle, L. R., & Cullers, D. K. 1996, *Icarus*, 119, 224
- Jordi, C., Gebran, M., Carrasco, J. M., et al. 2010, *A&A*, 523, A48
- Kocács, G., Zucker, S., & Mazeh, T. 2002, *A&A*, 391, 369
- Léger, A., Rouan, D., Schneider, J., et al. 2009, *A&A*, 506, 287
- Nutzman, P., & Charbonneau, D. 2008, *PASP*, 120, 317
- Pál, A., Bakos, G. Á., Torres, G., et al. 2008, *ApJ*, 680, 1450
- Tingley, B., & Sackett, P. D. 2005, *ApJ*, 627, 1011