

# Optimized transit detection algorithm to search for periodic transits of small planets

Michael Hippke<sup>1</sup> and René Heller<sup>2</sup>

<sup>1</sup> Sonneberg Observatory, Sternwartestr. 32, 96515 Sonneberg, Germany  
e-mail: michael@hippke.org

<sup>2</sup> Max Planck Institute for Solar System Research, Justus-von-Liebig-Weg 3, 37077 Göttingen, Germany  
e-mail: heller@mps.mpg.de

Received 16 November 2018 / Accepted 22 January 2019

## ABSTRACT

We present a new method to detect planetary transits from time-series photometry, the transit least squares (TLS) algorithm. TLS searches for transit-like features while taking the stellar limb darkening and planetary ingress and egress into account. We have optimized TLS for both signal detection efficiency (SDE) of small planets and computational speed. TLS analyses the entire, unbinned phase-folded light curve. We compensated for the higher computational load by (i.) using algorithms such as “Mergesort” (for the trial orbital phases) and by (ii.) restricting the trial transit durations to a smaller range that encompasses all known planets, and using stellar density priors where available. A typical K2 light curve, including 80 d of observations at a cadence of 30 min, can be searched with TLS in  $\sim 10$  s real time on a standard laptop computer, as fast as the widely used box least squares (BLS) algorithm. We perform a transit injection-retrieval experiment of Earth-sized planets around sun-like stars using synthetic light curves with 110 ppm white noise per 30 min cadence, corresponding to a photometrically quiet  $K_p = 12$  star observed with *Kepler*. We determine the SDE thresholds for both BLS and TLS to reach a false positive rate of 1% to be  $SDE = 7$  in both cases. The resulting true positive (or recovery) rates are  $\sim 93\%$  for TLS and  $\sim 76\%$  for BLS, implying more reliable detections with TLS. We also test TLS with the K2 light curve of the TRAPPIST-1 system and find six of seven Earth-sized planets using an iterative search for increasingly lower signal detection efficiency, the phase-folded transit of the seventh planet being affected by a stellar flare. TLS is more reliable than BLS in finding any kind of transiting planet but it is particularly suited for the detection of small planets in long time series from *Kepler*, TESS, and PLATO. We make our python implementation of TLS publicly available.

**Key words.** eclipses – methods: data analysis – planets and satellites: detection – planets and satellites: individual: TRAPPIST-1 – methods: statistical

## 1. Introduction

Since the first discovery of an extrasolar planetary transit across the disk of a distant star (Charbonneau et al. 2000), exoplanet surveys have expanded greatly in numbers and volume. Ground-based transit searches such as HATNet (Bakos et al. 2004), WASP (Pollacco et al. 2006), KELT (Pepper et al. 2007), and CHESPA (Zhang et al. 2019), and space-based search campaigns like CoRoT (Auvergne et al. 2009), *Kepler* (Borucki et al. 2010), K2 (Howell et al. 2014), and TESS (Ricker et al. 2014) produced vast data sets. These encompass hundreds of thousands of stars, cadences of seconds or minutes, and data sets that span several years. The PLATO space mission, with an expected launch in 2026 and a nominal six year duty cycle, will shadow these surveys by observing up to a million relatively bright stars with cadences of 25 s or 10 min (Rauer et al. 2014). These modern exoplanet transit searches require fast, sensitive, and reliable algorithms to detect the expected but unknown transit signals.

The box least squares (BLS) algorithm (Kovács et al. 2002, 2016) has become the standard tool for exoplanet transit searches in large data sets. It approximates the transit light curve as a (negative) boxcar function with a normalized average out-of-transit flux of zero and a fixed depth during the transit. This approach is key to its computational speed and allows for reliable detections of high to medium signal-to-noise ratio (S/N) signals, such as

large (Jupiter-sized) and moderately large (Neptune-sized) planets around sun-like stars in most surveys.

The BLS detection efficiency for low-S/N signals from Earth-sized planets around sun-like stars, however, is significantly smaller because the transit depths are comparable to the level of instrumental and stellar noise. Moreover, BLS introduces a systematic noise component that comes from the mathematical concept of the boxcar function. This binary model of a fixed out-of-transit and a fixed in-transit flux is equivalent to the neglect of the stellar limb darkening and of the planetary ingress and egress in the light curve. This box-shape approximation introduces an extra noise component in the test statistic that dilutes low-S/N signals. Here we present an improved transit search algorithm that attempts to minimize this systematic noise component in the search statistic.

The BLS algorithm has been analyzed and optimized in depth, for example in terms of the optimal frequency sampling, optimal phase sampling, and various other parameters (Ofir 2014). BLS has been extended to variable intervals between successive transits (Carter & Agol 2013), to work with non-Gaussian errors (Bouffleur et al. 2014), to improve speed at the cost of sensitivity (Renner et al. 2008), and to refine the detected transit parameters (Collier Cameron et al. 2006; Hartman & Bakos 2016). Further adaptations were made for the application to circumbinary planets (Ofir 2008).

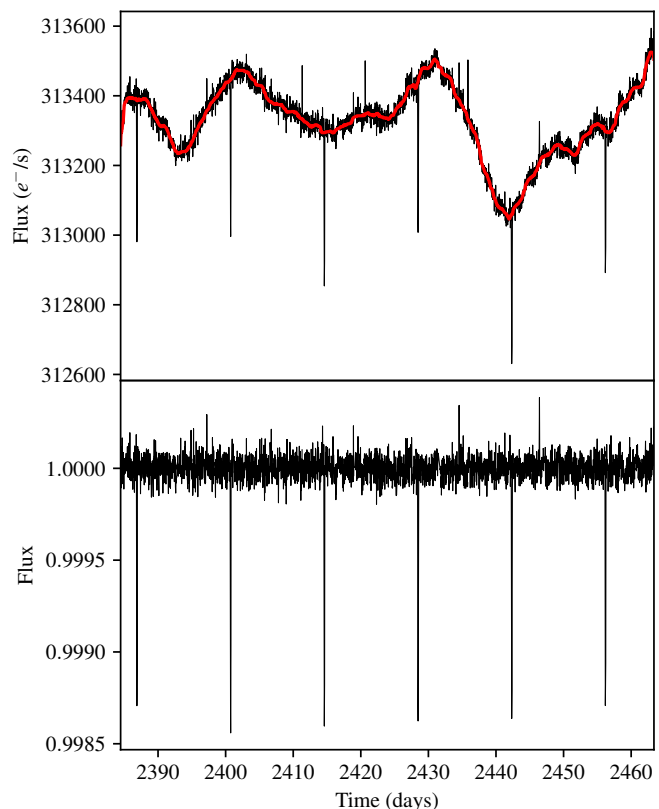
Alternatively, the “matched filter” algorithm is similar to BLS in modeling the transit as a boxcar, but it uses a different test statistic (Jenkins et al. 1996; Bordé et al. 2007). Phase dispersion minimization (Stellingwerf 1978) has been shown to be inferior to BLS for transit detection (Kovács et al. 2002). Analysis of variance (AoV, Schwarzenberg-Czerny & Beaulieu 2006) also uses a box-shaped transit model (which the authors refer to as top-hat), and has been demonstrated to have a lower detection efficiency than BLS in WASP data (Enoch et al. 2012). Bayesian algorithms can search for any signal form (Doyle et al. 2000; Defaÿ et al. 2001), but are not widely used. For example, the Gregory-Loredo method for Bayesian periodic signal detection uses step-functions (boxes with multiple steps) (Aigrain & Favata 2002; Aigrain & Irwin 2004). Wavelet-based algorithms (Régulo et al. 2007) are of similar detection efficiency and have been widely used for automated analyses of CoRoT (Régulo et al. 2009), *Kepler* (Jenkins et al. 2010), and TESS (Jenkins et al. 2016) data. Polynomials have also been suggested to approximate transit shapes more adequately than boxes (Cabrera et al. 2012) and this idea has been used (Johnson et al. 2016; Livingston et al. 2018), although without a comparison to BLS in terms of detection efficiency and computational effort.

Comparisons of different algorithms showed that BLS is the best of all known algorithms for weak signals (Tingley 2003a,b), but “no detector is clearly superior for all transit signal energies”, which has been verified by empirical tests of the methods (Moutou et al. 2005).

New techniques have now arrived with the advent of artificial intelligence. Deep learning algorithms are usually trained with a series of transit shapes (Pearson et al. 2018; Zucker & Giryes 2018; Armstrong et al. 2018). Random-Forest methods detect 7.5% more planets than classical BLS for low S/N transits (Mislis et al. 2016) because (many different) real transit shapes are used instead of a box. Disadvantages of these methods include substantial computational requirements, high implementation complexity, and a difficulty in understanding the origin of the results due to the many abstraction layers.

Here we present a new transit search algorithm that is easy to use, publicly available<sup>1</sup>, and has a detection statistic that is generally more sensitive than that of BLS. Most important, it is optimized to find small planets in large data sets. The algorithm assumes a realistic transit shape with ingress and egress and stellar limb darkening (as per Mandel & Agol 2002) using a predefined parameterization that we optimized based on all previous exoplanet transit detections. The resulting increase in the detection significance of the algorithm by 5–10% comes at the toll of larger computational demands. Given the tremendous growth of available CPU power in the past 60 years (Moore 1965) and in particular since the publication of the BLS algorithm in 2002, however, we argue that CPU margins are not as crucial to the detection of small planets as is the significance of the test statistic. That said, we have nevertheless optimized the algorithm for computational speed as far as possible.

Our algorithm is particularly suited for the detection of Earth-sized planets with *Kepler*/K2, TESS, or with the future big data sets from the PLATO mission. In fact, the improvements of our new algorithm are most substantial for small planets with few transits, which is a common characteristic of Earth-sized planets in the habitable zones around sun-like stars.



**Fig. 1.** *Top panel:* K2 long cadence light curve of the star K2-110, which exhibits transits of a mini-Neptune-sized planet, K2-110 b. The black line shows the light curve that has been corrected for instrumental effects with EVEREST and the red line shows our running median of 51 data points. *Bottom panel:* EVEREST light curve divided by the running median.

## 2. Methods

We illustrate the methodology of TLS using the K2 light curve of the metal-poor K3 dwarf star K2-110, which hosts a transiting massive mini-Neptune (K2-110 b, EPIC 212521166b) in a 13.86 d orbit (Osborn et al. 2017). Huber et al. (2016) estimate a stellar effective temperature of  $T_{\text{eff}} = 4628$  K, a surface gravity of  $\log(g) = 4.6$ , a stellar mass of  $M_s = 0.752 M_{\odot}$ , and a stellar radius of  $R_s = 0.7 R_{\odot}$ .

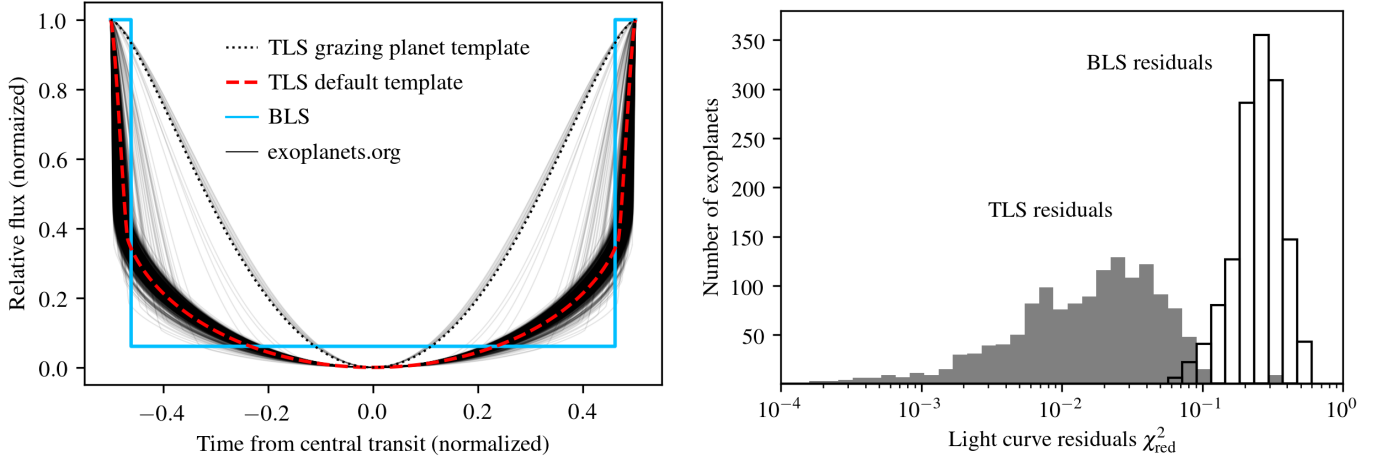
In the top panel of Fig. 1 we show the K2 light curve after correction for instrumental effects with EVEREST (Luger et al. 2016) (black line) together with the running median of 51 exposures (red line). The lower panel of Fig. 1 shows the light curve after division by the running median. This is the data used throughout this section and we note that this pre-processing or detrending of the light curve is not part of TLS.

Readers interested in detrending techniques are referred to the Savitzky & Golay (1964) filter (used by Gilliland et al. 2011), the median filter (Tal-Or et al. 2013), polynomial filters (Gautier et al. 2012; Rodenbeck et al. 2018), the Cosine Filtering with Autocorrelation Minimization (CoFiAM, Kipping et al. 2013), and Gaussian process (Aigrain et al. 2015).

### 2.1. Transit model

The key idea of TLS is to search for transits using a transit-like search function rather than a box. Our first task then is to identify the transit light curve that is most representative of the known exoplanet transit light curves, assuming that future exoplanet

<sup>1</sup> <http://github.com/hippke/tls>



**Fig. 2.** *Left panel:* transit light curves for all 2346 transiting *Kepler* planets from the Exoplanet Orbit Database (as of 1 November 2018) with  $R_p/R_s < 0.05 \sim 5.4 R_{\oplus}/R_{\odot}$  (black lines, one per planet). The default template for TLS is shown with a red dashed line and the optional grazing planet template is shown with a black dotted line. The best-fitting box is shown as a light blue solid line for comparison. *Right panel,* gray histogram: reduced  $\chi^2$  residuals between the TLS default (median) transit template and the real transit light curves in the left panel. Open histogram: reduced  $\chi^2$  residuals between the box and the real transit light curves. The separation between the two histograms confirms that the TLS default transit template is a substantially better match to the observations than a box in  $\geq 99\%$  of the cases.

detections will be done most efficiently with this particular function. We refer to this function as the default template for TLS and describe its construction in the following. Although this decision of using a transit curve template to match previous detections might suggest that TLS will inherit a detection bias and search mostly for planets like the ones we already know, we have verified that TLS is better than BLS in finding any kind of planet, in particular grazing transmitters and very small planets.

In fact, we decide to optimize the TLS template for the detection of small planets. We have verified that large planets, which produce deep transits, can also be found with this TLS template with a higher signal detection efficiency than with a box. In turn, if we had chosen to optimize the template to find large planets, then small planets would be more likely to be missed. We impose an arbitrary limit of  $R_p/R_s < 0.05 \sim 5.4 R_{\oplus}/R_{\odot}$  on the planet-to-star radius ratio and retrieve the orbital inclination ( $i$ ), semimajor axis in units of stellar radii ( $a/R_s$ ),  $R_p/R_s$ , and the orbital period ( $P$ ) for all transiting *Kepler* planets from the Exoplanet Orbit Database<sup>2</sup> (Wright et al. 2011). We set the orbital eccentricity of each planet to zero and obtain the predicted limb darkening coefficients  $c_1$  and  $c_2$  of a quadratic limb darkening law for each host star, using the stellar effective temperature ( $T_{\text{eff}}$ ) and surface gravity ( $\log(g)$ ), from the stellar model atmospheres of Claret et al. (2012, 2013) in the *Kepler* bandpass.

In the left panel of Fig. 2 we plot the 2346 resulting model transit light curves, normalized to the transit depth and transit duration, using the *batman* implementation (Kreidberg 2015a,b) of the Mandel & Agol (2002) analytic transit model with quadratic limb darkening (black lines, one for each planet). We then construct the TLS default template transit curve from the median values of the above-mentioned input parameters to the analytic transit model for quadratic limb darkening (red dashed line). We emphasize that the TLS template is not a fit to the observed normalized transit light curves in Fig. 2 but a model light curve based on the median input parameters of all known transiting exoplanets. In fact, the user of TLS is free to choose their own parameterization of a template transit light curve for their search. We also show our template for a grazing transiting planet (black dotted line) and the best fitting box.

<sup>2</sup> <http://exoplanets.org>

In the right panel of Fig. 2 we show the reduced  $\chi^2$  residuals between the TLS template and the real transit light curves (gray histogram) and the reduced  $\chi^2$  residuals between the BLS box and the real transit light curves (open histogram). The two histograms show a substantial offset with the TLS template resulting in much smaller  $\chi^2$  residuals than the box used for the BLS algorithm. We note that the single outlier of the TLS distribution at  $\chi_{\text{red}}^2 \sim 0.35$  belongs to KOI-7880, which is a grazing transit planet.

## 2.2. Transit search statistic

The TLS algorithm searches for periodic transit-shaped signals in time series of flux measurements. The algorithm operates by phase-folding the data over a range of trial periods ( $P$ ), transit epochs ( $t_0$ ), and transit durations ( $d$ ). It then calculates the  $\chi^2$  statistic of the phase-folded light curve between the  $N$  data points of the respective transit model ( $y_i^m$ ) and the observed values ( $y_i^o$ ) as per

$$\chi^2(P, t_0, d) = \sum_{i=1}^N \frac{(y_i^m(P, t_0, d) - y_i^o)^2}{\sigma_i^2}, \quad (1)$$

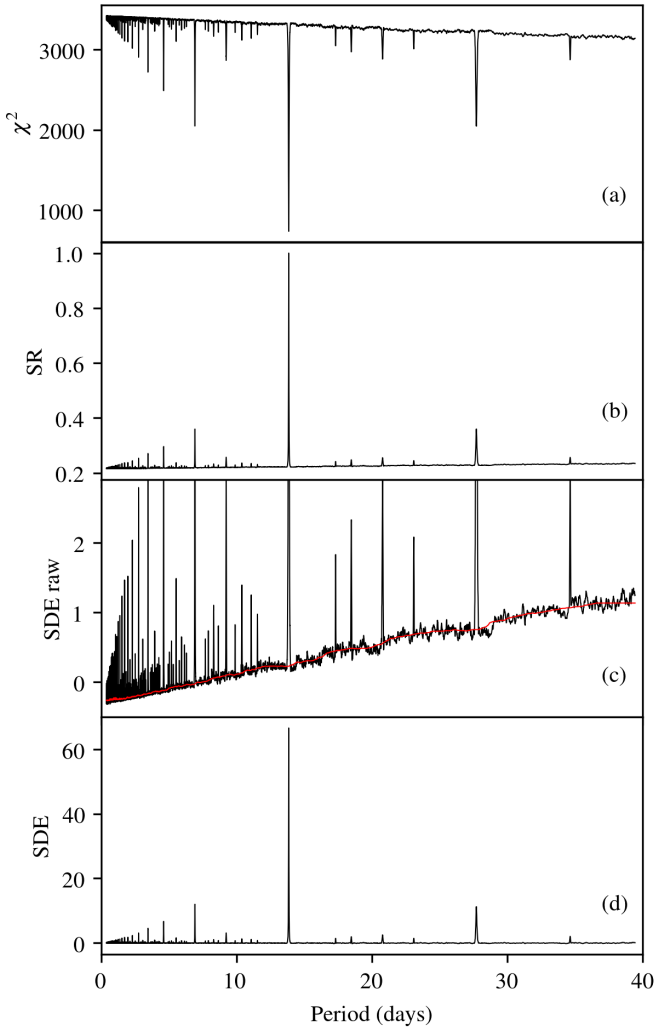
where  $\sigma_i^2$  are the standard deviations in the light curve. In Fig. 3a, we show the spectrum of minimum  $\chi^2$  as a function of  $P$ . In other words, for each trial period TLS searches the minimum  $\chi^2$  by testing all combinations of the  $(t_0, d)$  2D hyperspace of the 3D parameter grid. TLS uses the global  $\chi^2$  minimum

$$\chi_{\text{min, glob}}^2 \equiv \min(\chi^2(P, t_0, d)) \equiv \chi^2(P', t'_0, d'), \quad (2)$$

at the location  $(P', t'_0, d')$  of our 3D parameter space for the normalization of the test statistic below. In Fig. 3a, we locate  $\chi_{\text{min, glob}}^2$  at about 13.87 d, corresponding to the published value of K2-110 b by Osborn et al. (2017).

TLS uses a modified version of the test statistic originally implemented in the BLS algorithm<sup>3</sup>, the signal detection efficiency (SDE; Alcock et al. 2000). The SDE has been widely

<sup>3</sup> In Appendix B we identify a glitch in a patch to the previously known BLS edge effect that slightly affects the test statistic. This has been corrected for in TLS.



**Fig. 3.** *Panel a:* distribution of  $\chi^2$  (minimized over  $t_0, d$ ) obtained by phase-folding the light curve of K2-110 b over different trial periods. *Panel b:* signal residue for the best fitting periods throughout the parameter space. *Panel c:* raw signal detection efficiency (black line) and walking median (red line). *Panel d:* signal detection efficiency used by TLS. This is the result of the division of the raw SDE by its walking mean in *panel c*.

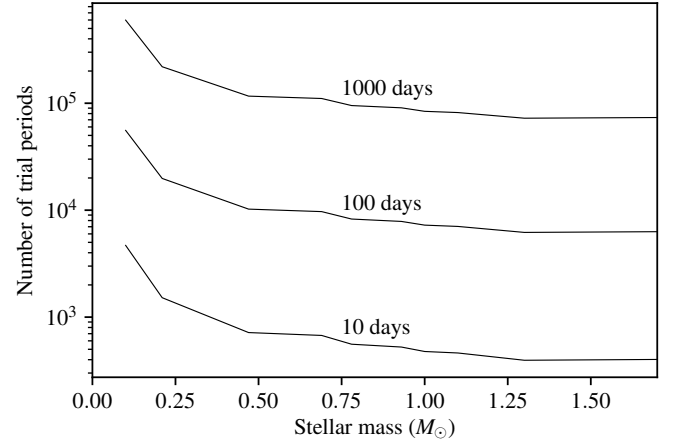
demonstrated to yield useful results and it has become a standard metric in the exoplanet hunting community. Our implementation of TLS, however, does not apply a binning of the phase-folded light curve to compute the signal residue (SR) between the model and the data. Our approach is computationally more expensive (quadratic in the number of data points, see Sect. 3.4) but key to making TLS more sensitive to the signals of small transiting planets.

TLS calculates the SR from the distribution of minimum  $\chi^2$  as a function of  $P$ ,

$$\text{SR}(P) = \frac{\left(\frac{1}{\chi_{\min}^2(P)}\right)}{\left(\frac{1}{\chi_{\min, \text{glob}}^2}\right)} = \frac{\chi_{\min, \text{glob}}^2}{\chi_{\min}^2(P)}, \quad (3)$$

which necessarily results in the  $\text{SR}(P)$  distribution to range between 0 and 1 (see Fig. 3b).

The  $\text{SDE}(P)$  distribution is then obtained as per Kovács et al. (2002) using the arithmetic mean  $\langle \text{SR}(P) \rangle$ , the standard deviation



**Fig. 4.** Optimal number of trial periods ( $N_{\text{freq, opt}}$ ) as a function of stellar mass for three different time spans of a hypothetical stellar light curve.

$\sigma(\text{SR}(P))$ , and the peak value  $\text{SR}_{\text{peak}}$  of  $\text{SR}(P)$  via

$$\text{SDE}(P) = \frac{\text{SR}_{\text{peak}} - \langle \text{SR}(P) \rangle}{\sigma(\text{SR}(P))}. \quad (4)$$

With  $\text{SR}_{\text{peak}} = 1$  by definition in Eq. (3), we have

$$\text{SDE}(P) = \frac{1 - \langle \text{SR}(P) \rangle}{\sigma(\text{SR}(P))}. \quad (5)$$

An SDE value of  $x$  for any given  $P$  means that the statistical significance of this period is  $x\sigma$  compared to the mean significance of all other periods. We refer to the resulting  $\text{SDE}(P)$  distribution in Fig. 3c as the raw SDE. The final step in our construction of the transit search statistic is in the removal of the systematic noise component that is inherent to the SDE distribution as explained by Ofir (2014). We follow this author in removing this trend with a walking median filter through the  $\text{SDE}(P)$  periodogram, the result of which is shown in Fig. 3d for K2-110. The transit signal of Kepler-110 b can be found at 13.8662 with an SDE value of 66.7. We determine a conservative error estimate via the half width at half maximum of the SDE peak, which is 0.0122 d. Hence, TLS determines the orbital period of K2-110 b as  $P = 13.8662 (\pm 0.0122)$  d.

Empirical thresholds for a transit detection have been proposed to range from  $\text{SDE} > 6$  (Dressing & Charbonneau 2015),  $\text{SDE} > 6.5$  (Livingston et al. 2018),  $\text{SDE} > 7$  (Siverd et al. 2012), 6–8 as a function of period (Pope et al. 2016; Aigrain et al. 2016), and up to 10 (Wells et al. 2018). Lower SDE thresholds mean better completeness but also higher false alarm rates.

### 2.3. The TLS parameter grid

#### 2.3.1. Period sampling

In the search for sine-like signals, for example using Fourier transforms, a uniform sampling of the trial frequencies is usually quite efficient. A uniform sampling of the orbital frequencies has also been suggested for BLS (Kovács et al. 2002). As shown by Ofir (2014), however, this sampling of the orbital frequency tends to be insensitive to either short- or long-period planets, but it is always computationally inefficient. Ofir (2014) derived the optimal number of test frequencies as

$$N_{\text{freq, opt}} = \left( f_{\max}^{1/3} - f_{\min}^{1/3} + \frac{A}{3} \right) \frac{3}{A}, \quad (6)$$

with

$$A = \frac{(2\pi)^{2/3}}{\pi} \frac{R_s}{(GM_s)^{1/3}} \frac{1}{S \times OS}, \quad (7)$$

where  $G$  is the gravitational constant,  $S$  is the time span of the data set, and  $OS$  is the oversampling parameter to be chosen between 2 and 5 to ensure that the SDE peak is not missed between trial frequencies (or periods). The minimum and maximum trial orbital frequencies can be found at  $f_{\min} = 2/S$  (or  $f_{\min} = 3/S$  if three transits are required) and at the most short-period (high-frequency) circumstellar orbit, the Roche limit,  $f_{\max} = \sqrt{GM_s}/(3R_s)^3/(2\pi)$ . Strictly speaking, the Roche limit depends on the density ( $\rho_p$ ) of the planet, and the term  $3R_s$  for our expression of  $f_{\max}$  assumes the most pessimistic case of an extremely low-density fluid-like planet with  $\rho_p = 1 \text{ g cm}^{-3}$ , which can be compared to Jupiter's mean density of  $1.33 \text{ g cm}^{-3}$ . Our TLS implementation generates an array of evenly spaced orbital frequencies with  $N_{\text{freq,opt}}$  constant steps between  $f_{\min}$  and  $f_{\max}$  and then computes the (non-uniform) trial orbital periods as the inverse of this frequency grid.

Since computational speed is a key concern for us, we illustrate the resulting number of trial periods in Fig. 4, using three different time spans  $S$  of a hypothetical light curve for different stellar masses (and radii, assuming main-sequence mass-radius relation). We find that an extension of the light curve by a certain factor – here ten between the three example curves – increases the number of trial periods by the same factor. This plot warns us of the large number of trial periods that need to be examined for planets around low-mass stars, with  $N_{\text{freq,opt}}$  reaching values of up to almost one million for a light curve with 1000 d of continuous observations of a very-low-mass star. This feature is inherent to both TLS and BLS.

### 2.3.2. Transit depth

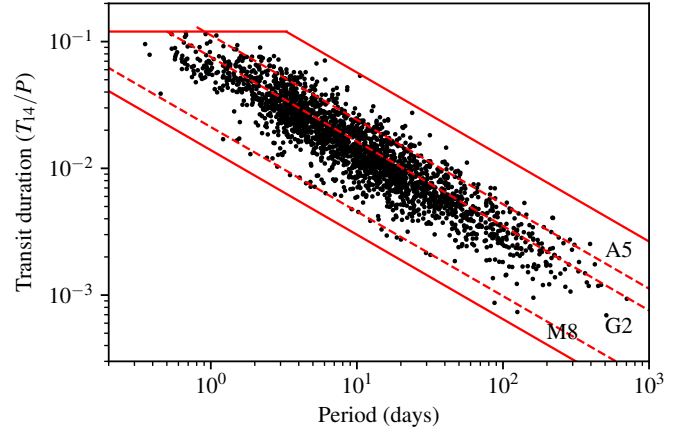
TLS measures the mean flux of the in-transit data points under consideration. It then calculates the corresponding maximum transit depth  $\delta$  and the resulting planet-to-star radius ratio under the assumption of zero transit impact parameter using the analytic solutions for common stellar limb darkening laws found by Heller (2019). With this calculation, the signal shape is scaled and compared to the data points.

TLS can also be used to search for user-defined signal shapes (for example, flares), either with positive or negative flux. If an analytical scaling option is not available, TLS can perform a numerical iterative fit using an initial guess based on the mean of the in-signal data,  $\langle y_{i,\text{in}}^o \rangle$ . The  $\delta$  range to be tested with TLS is bracketed between  $\langle y_{i,\text{in}}^o \rangle/10$  and  $10 \times \langle y_{i,\text{in}}^o \rangle$ . TLS uses an iterative ternary algorithm (Knuth 1998) to tighten the interval in steps of 1/3 until the upper and lower limits differ by  $<X\%$  in signal depth (or height), where  $X$  is a user-defined threshold.

### 2.3.3. Transit duration

Transit searches using the BLS algorithm usually operate with trial transit durations<sup>4</sup>  $T_{14}$  that span 0.00125–0.07 (Petigura et al. 2013), 0.01–0.1 (Giacobbe et al. 2012), or 0.001–0.2 (Sanchis-Ojeda et al. 2014; Aigrain et al. 2016) times the orbital period. More than half of the corresponding  $T_{14}(P)$

<sup>4</sup> We follow the common nomenclature to indicate the time interval between the first and fourth contact of the stellar and planetary silhouettes as  $T_{14}$  (Seager & Mallén-Ornelas 2003).



**Fig. 5.** Transiting planets from the Exoplanet Orbit Database in the  $T_{14}/P$ - $P$  diagram. BLS implementations typically search a linearly spaced uniform grid, or roughly the entire diagram. However, more than half of this search space is not populated with planets. The default parameterization of TLS only searches inside the area embraced by the solid lines, which are defined in Eq. (10). TLS users can nevertheless redefine their own cuts of the  $T_{14}/P$ - $P$  diagram and search for planets with hitherto unknown properties in this diagram.

diagram, however, is not populated with exoplanet discoveries in these regions of the parameter space (see Fig. 5). We explain this absence of transiting planets using geometrical constraints and *Kepler's* third law.

For example, there are no planets known with  $P = 10 \text{ d}$  and  $T_{14}/P < 5 \times 10^{-3}$ . From an astrophysical perspective, only extremely eccentric planets could have such a short transit duration – apparently a very rare, or even non-existent kind of exoplanet. It thus appears reasonable to us to restrict the computational effort to the physically plausible regions of the  $T_{14}/P$  diagram. We also conclude from Fig. 5 that the transit duration search grid shall be linear in log-space.

For circular orbits, the maximum transit duration is  $T_{14,\text{max}} = 2(R_s + R_p)/v_p$ , where  $v_p$  is the planet's average orbital velocity during the transit. Shorter transit durations are possible if the planetary transit path is not across the stellar diameter. We then have

$$T_{14,\text{max}} = \frac{2(R_s + R_p)}{v_p} = \frac{2(R_s + R_p)}{2\pi a/P}. \quad (8)$$

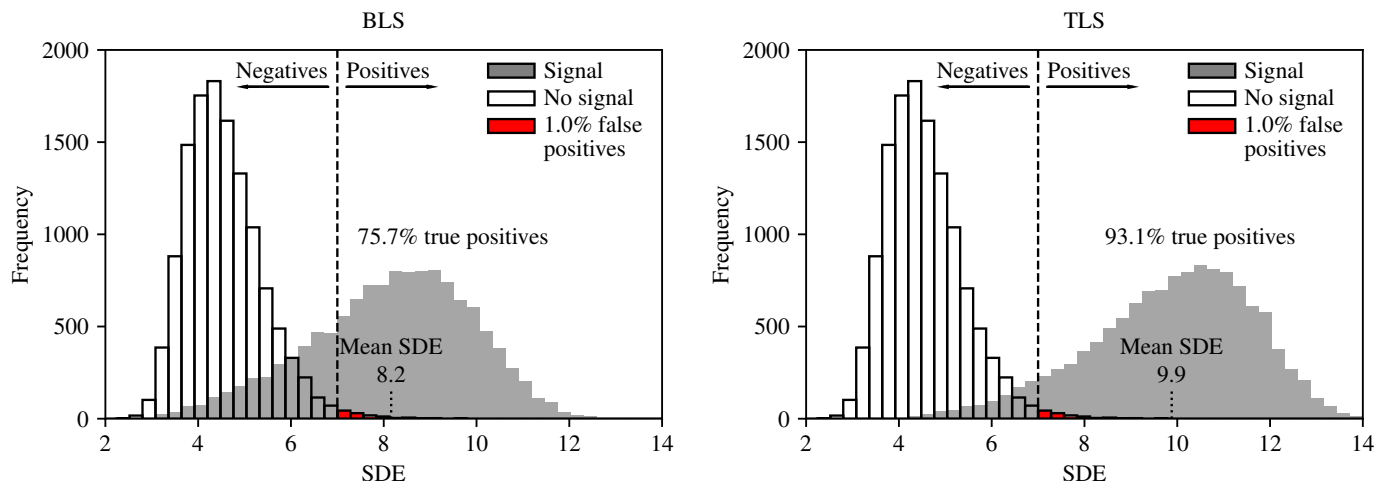
In the limit of the star being much more massive than the planet, *Kepler's* third law becomes

$$\left(\frac{2\pi}{P}\right)^2 a^3 \approx GM_s \Leftrightarrow a \approx \left(GM_s \left(\frac{P}{2\pi}\right)^2\right)^{1/3}. \quad (9)$$

We insert Eq. (9) into Eq. (8) and obtain

$$\begin{aligned} T_{14,\text{max}} &= \frac{(R_s + R_p)P}{\pi} \left(\frac{1}{GM_s} \left(\frac{2\pi}{P}\right)^2\right)^{1/3} \\ &= (R_s + R_p) \left(\frac{4P}{\pi GM_s}\right)^{1/3}. \end{aligned} \quad (10)$$

In Fig. 5 we plot Eq. (10) for  $R_p = 2R_J$  ( $R_J$  being Jupiter's radius) orbiting an A5 star ( $M_s = 2.1 M_\odot$ ,  $R_s = 1.7 R_\odot$ ) to maximize the effect of a very large planet on the transit duration. To



**Fig. 6.** Statistics of the signal detection efficiency for a transit injection-retrieval experiment of simulated light curves with white noise only. *Left panel:* box least squares algorithm. *Right panel:* transit least squares algorithm. Both panels show the results of 10 000 realizations of a 3 yr light curve with white noise only (open histograms) and of the same amount of light curves with white noise and an Earth-like planetary transit around a G2V star (gray histograms). Outlined histograms relate to the SDE maximum value in a noise-only search. Gray histograms refer to the highest SDE value within 1% of the period of the injected transit. The SDE thresholds at which the false positive rates are 1% is found to be  $SDE_{fp=1\%} = 7$ . At this SDE threshold, the recovery rate of the injected signals (the true positive rate) is 75.7% for BLS and 93.1% for TLS, while the fraction of missed signals (the false negative rate) is 24.3% for BLS and 6.9% for TLS.

embrace the physically plausible search space, we also show a main-sequence M8 red dwarf star ( $M_s = 0.1 M_\odot$ ,  $R_s = 0.13 R_\odot$ ) and a sun-like star, both with a small planet ( $R_p = R_\oplus$ ). We see that a significant amount of planets is actually located above the uppermost of these lines, which can be attributed to bloated super-Jovian planets and/or to planets transiting slightly evolved stars, for example. The absence of planets in the lower left part of the diagram could be partly astrophysical in nature and interpreted as a trace of planet formation and evolution (for example, the absence of ultra-short period planets around M dwarfs). That said, planets can naturally have transit durations that are arbitrarily shorter than  $T_{14,max}$ , for example, on eccentric or inclined orbits. The empty space can nevertheless be explained with a detection bias against planets with transit durations of just a few minutes, for example,  $\sim 15$  min for  $P = 1$  d and  $T_{14}/P = 10^{-2}$ .

In order to compensate for planets transiting evolved stars as well as for planets on eccentric orbits<sup>5</sup> and other astrophysical effects that are potentially hard to predict, we use Fig. 5 to derive empirical estimates for the maximum and minimum values of  $d$  to be searched. We parameterize the upper limit via Eq. (10) and using  $M_s = 1 M_\odot$  and  $R_s = 3.5 R_\odot$ , and we parameterize the lower limit of  $T_{14}$  using  $M_s = 1 M_\odot$  and  $R_s = 0.184 R_\odot$  in Eq. (10). We note that these two parameterizations do not correspond to any particular or even physically plausible main-sequence star. The motivation behind this parameterization of Eq. (10) is entirely empirical with the aim of embracing all known transiting exoplanets. Searches for planets around more exotic stars, such as white dwarfs, require other limits. Using our TLS implementation, the user can conveniently set arbitrary limits of their choice.

Our default empirical limits for the transit durations to be searched with TLS are shown with inclined solid lines in Fig. 5 and their parameterization was intentionally chosen to encompass all known transiting exoplanets. The horizontal cutoff at

$T_{14}/P = 1.12 \times 10^{-1}$  is a global threshold. This is the default region in the  $T_{14}$ -( $P$ ) diagram to be tested with our implementation of the TLS algorithm. That said, the user can define their own range of transit durations to be tested.

### 3. Results

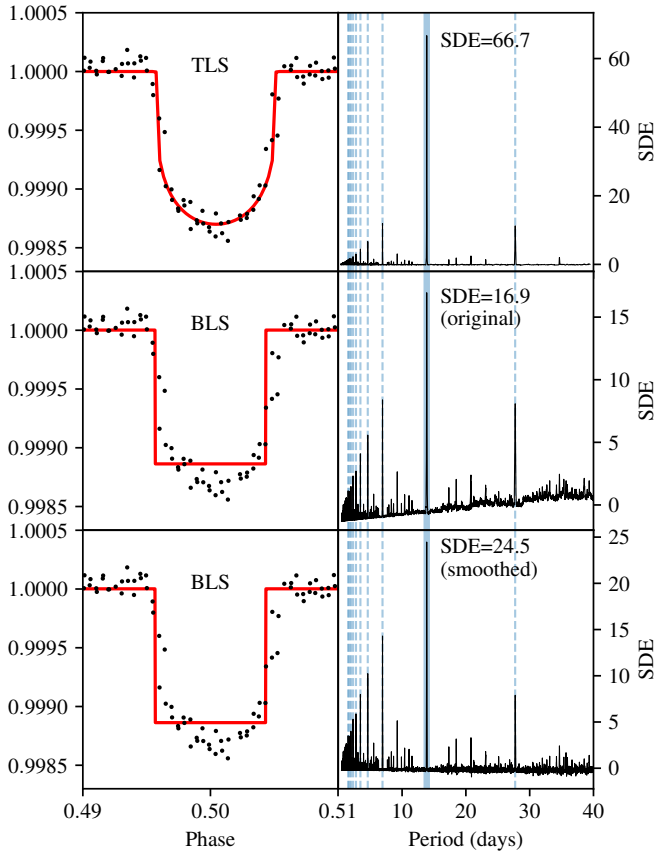
#### 3.1. TLS and BLS signal detection efficiency for white noise

As a first test of the performance of TLS in comparison to BLS we generated synthetic light curves with white noise only. These light curves have a time span of  $S = 3$  yr and a cadence of 30 min with 110 ppm noise (standard deviation) per cadence. This noise level is adapted to a best-case scenario for *Kepler* data, where the total noise over 6.5 h was found to be 30 ppm for a  $K_p = 12$  star (Gilliland et al. 2015). We note that the noise in the *Kepler* light curves does not have Gaussian properties due to instrumental and stellar trends (for example, from stellar rotation), all of which complicates transit detections in practice.

We injected three transits of an Earth-sized planet around a sun-like star with transit impact parameters randomly chosen in  $b = [0, 1]$  and with solar quadratic limb darkening as seen in the *Kepler* bandpass into a set of 10 000 pre-computed synthetic light curves with different white noise realizations. Then we conducted both a TLS search and a standard BLS search, in both of which we used our optimized period grid (Sect. 2.3.1) for a fair comparison. A detection was counted as “positive” if the highest peak in the power spectrum was within 1% of the injected transit period.

We determine the SDE thresholds for a false positive rate of 1% to be  $SDE_{fp=1\%} = 7$  for BLS and TLS. Given these thresholds and the (forced) false positive rates of 1%, TLS recovers 93.1% of the injected signals (the true positives) compared to 75.7% for BLS. While the SDE distribution of the white noise-only light curves is virtually identical for both BLS and TLS, the SDE distribution of the light curves with signal is shifted to higher SDE values for TLS, with a mean of  $\langle SDE \rangle_{TLS,p} = 9.9$  compared to a mean of  $\langle SDE \rangle_{BLS,p} = 8.2$  for BLS. In Fig. 6, we illustrate the results of our experiment.

<sup>5</sup> In eccentric cases the average orbital velocity in-transit can be smaller or larger and the resulting transit duration can be larger or smaller than in the circular case.



**Fig. 7.** Phase-folded transits and peridograms of *Kepler* K2-110b with TLS (*top*) and BLS (*middle, bottom*) fitting for the same trial periods and durations. We note the boost in signal detection efficiency from 16.9 with the original BLS, or 24.5 with the median-smoothed BLS to 64.2 with TLS. The vertical dashed blue lines denote the aliases of the period detected at the highest SDE value, respectively.

### 3.2. Comparison of TLS and BLS for K2-110b

We now compare the performance of TLS and BLS using a known planet with a high signal-to-noise ratio, K2-110b (EPIC 212521166), a massive mini-Neptune orbiting an old, metal-poor K3 dwarf star with a 13.9d period (Osborn et al. 2017). We retrieve estimates for the stellar mass ( $0.752^{+0.053}_{-0.044} M_{\odot}$ ), radius ( $0.7^{+0.048}_{-0.045} R_{\odot}$ ), effective temperature ( $T_{\text{eff}} = 4841^{+174}_{-145}$  K), and surface gravity ( $\log(g) = 4.628^{+0.025}_{-0.03}$ ) from the *Kepler* K2 EPIC catalog (Huber et al. 2016) using the automated `catalog_info` function of TLS. `catalog_info` also retrieves the quadratic limb-darkening coefficients for the *Kepler* bandpass ( $a = 0.7010$ ,  $b = 0.0462$ ) via a cross-match of the Claret et al. (2012) tables based on  $T_{\text{eff}}$  and  $\log(g)$ . With these priors, TLS creates an optimal period grid with an oversampling factor of five, which results in 21 500 trial periods between  $P_{\text{min}} = 0.4$  d and  $P_{\text{max}} = d/2 = 40$  d.

We run BLS and TLS searches with the same period grid, oversampling, and duration constraints (Fig. 7). TLS delivers a much higher SDE (66.7) compared to BLS (16.9, or 24.5 when median-smoothed). This is despite the fact that our priors are slightly different to the improved posteriors from Osborn et al. (2017), which suggest a hotter star ( $T_{\text{eff}} = 5050 \pm 50$  K) with the same surface gravity, resulting in different quadratic limb darkening parameters ( $a = 0.5322$ ,  $b = 0.1787$ ).

### 3.3. Recovery of the TRAPPIST-1 planets

Moving on to real light curves of Earth-sized planets with time-correlated (red) noise components, we chose the K2 light curve of TRAPPIST-1 as a stress-test for TLS to ensure its robustness. The system exhibits noise from instrumentals, stellar rotation, flares, and other sources, which can only imperfectly be removed using detrending. Signals of multiple planets occur with overlapping transits. Each planet produces its own set of harmonics and subharmonics in the power spectrum. To be considered robust, a detection algorithm must be able to handle these difficulties and our attempt of recovering this previously reported series of transit signals in retrospect offers an exquisite case to simulate the TLS search for Earth-sized planets.

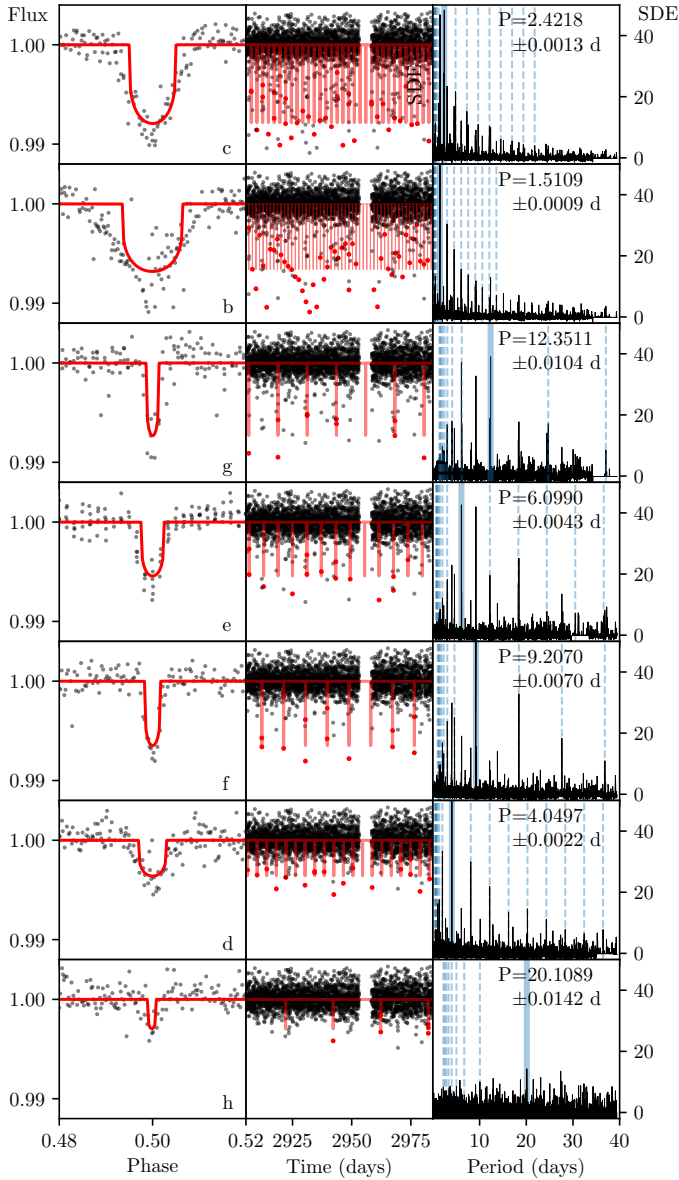
The well-studied TRAPPIST-1 system exhibits transits of seven terrestrial planets (Gillon et al. 2016, 2017) in a resonant chain, where the orbital periods are near-ratios of small integers (Luger et al. 2017). An automatic recovery of all planets is certainly difficult because of the low S/N of the individual transits resulting from the dim host star, the very small (Earth-sized) planets, transit timing variations, stellar flares, systematic trends from the stellar rotation of 3.3 days, and overlapping multi-planet transits.

As before, we use the K2 EVEREST data spanning 79 days in campaign 12. We divide the data by a running median of 13 data points.

On the one hand, the length of the running median filter window must be larger than the transit duration to prevent the transit signal from being distorted by the median filter. On the other hand, the window length must be sufficiently short to remove stellar variability. For planets around TRAPPIST-1, the longest plausible transit duration at a period of 40 days is 1.6 h (3 cadences). Our choice of 13 cadences is  $\sim 4$  times longer than the critical value. Longer median filter windows increase the residuals of the stellar noise significantly due to the strong variability of TRAPPIST-1. For comparison, in the case of K2-110b (Sect. 3.2), the amplitude of the stellar variability is much smaller and it occurs on timescales that are much longer than the transit duration. Therefore, we set the window length to 51 points ( $\sim 2$  d), but a window length of 1 d yields virtually the same result.

We then remove data points that deviate positively from the mean flux by more than  $+3\sigma$  to eliminate bright flares. A detailed analysis by Ducrot et al. (2018) carefully identifies transits affected by flares, incomplete transits, and multi-planet transits, which increases the quality of the in-transit data. Such a fine-tuned processing is beyond the scope of our analysis.

For our TLS search, we use the default template and priors on the stellar mass ( $0.089 \pm 0.006 M_{\odot}$ ) and radius ( $0.121 \pm 0.003 R_{\odot}$ ) with limb darkening for the *Kepler* bandpass of an M8-star (Van Grootel et al. 2018). The first run results in an SDE  $\sim 45$  detection of a signal with a period of  $P = 2.4218 \pm 0.0013$  d, which we identify as planet “c” (Fig. 8). Then we mask the in-transit data points of this signal using the TLS convenience function `transit_mask` and re-run TLS iteratively. Each successive run results in the detection (and masking) of planets c–b–g–e–f–d. The order of detection is based on the signal-to-noise ratio of the stacked transits. While planet “b” nominally has the highest S/N, its transit shape differs significantly from the TLS default template, making planet “c” have the highest SDE in the first TLS run. The seventh and outermost planet “h” is not automatically detected by TLS. We attribute this to the low number of transits (four) and to a flare that happened during the fourth transit, resulting in several in-transit data points showing unusually

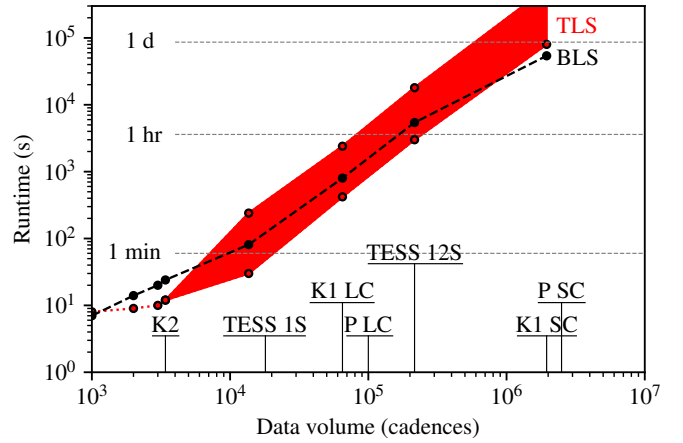


**Fig. 8.** Demonstration of the TLS performance on the TRAPPIST-1 system. *Left panels:* phase-folded transit light curve for the respective period and epoch at SDE maximum (black dots). The best-fit transit model (fitted for transit duration and depth) with quadratic stellar limb darkening is shown with a red solid line. Planet names are indicated in the lower right corner of each panel. Planets are sorted from top to bottom in the order of detection from an iterative TLS search of the K2 light curve. Planet “h” (*bottom panels*) is a false positive and not related to the actual detection of TRAPPIST-1 h (see Sect. 3.3). *Middle panels:* entire K2 light curve of TRAPPIST-1 with the detected in-transit data points highlighted in red. Transits detected in previous iterations were masked. *Right panels:* SDE( $P$ ) diagram for the light curve shown in the center.

high flux. We show the highest SDE peak (caused by noise) of this last search in the bottom row of Fig. 8 together with the best-fit transit shape, which is very noisy. This false positive signal illustrates the limits of automatic planet recovery, which apply to both TLS and BLS. We also verified that BLS is not able to detect planet “h” using the same data processing.

### 3.4. Computational costs

TLS aims at maximizing sensitivity while our implementation of TLS aims at maximizing computational speed at the same time.



**Fig. 9.** Algorithm run times for different missions, durations, and cadences. *Kepler* LC (30 min) and SC (1 min) are shown for 4.25 yr worth of data, *K2* assumes 80 d of LC data. *TESS* is represented with one (1S) and twelve (12S) seasons at 2 min cadence, respectively. *PLATO* light curves are considered over 2 yr in both 25 s short cadence (P SC) and 10 min long cadence (P LC). The red area shows the full range of TLS run times. The upper end assumes no priors on stellar density and fitting signals down to 10 ppm. The lower end assumes typical priors from catalog data and a 100 ppm threshold (or a 1% threshold on phase sampling). In the latter case, run times are shorter than BLS for all but the largest data sets. We note the slope of roughly two orders of magnitude of the run time per order of magnitude of cadences.

Since computing power has been continuously increasing for more than half a century now (Moore 1965) and since the whole point of TLS is to offer unprecedented sensitivity, we prioritized the latter over computational efficiency whenever necessary (for example: no phase binning).

The computational effort per light curve is a complicated function of the stellar mass and stellar radius (see Fig. 4) and it depends linearly on the time span of the light curve ( $S$ ). Another important factor is the number of trial epochs (or trial phases), which increases linearly with  $S$  for a constant cadence. Both things combined, we find that the computational load increases quadratically with  $S$ , which itself is proportional to the number of data points for a fixed cadence.

The default TLS configuration has a typical run time per *Kepler* K2 long cadence (K2 LC) light curve ( $\sim 4000$  data points,  $S = 80$  d) of 10 s, virtually identical to BLS on the same Intel Core i7-7700K (Fig. 9). We used the reference implementation of BLS provided by Astropy 3.1 (Astropy Collaboration 2013, 2018) in the C programming language, parallelized with the OpenMP interface<sup>6</sup>. To compare K2 run times, we used the same number of trial transit durations (66) in both algorithms, and the same optimal grid of  $\sim 10\,000$  periods as determined by TLS. We note that the optimal grid is not available by default in Astropy and most other BLS implementations, and only used by part of the community, resulting in three to five times longer run times for BLS at the same level of sensitivity. TLS run times are strongly dependent on the stellar density prior and the shallowest transit depth considered for fitting. A range of plausible values for optimistic and pessimistic cases is shown in Fig. 9 and explained in more detail in Appendix A.

Our measurements are in agreement with “one minute run time [for BLS] per processor core per K2 campaign star” (Vanderburg et al. 2016). Some years ago, BLS performance

<sup>6</sup> <http://docs.astropy.org/en/stable/api/astropy.stats.BoxLeastSquares.html>

was noted as “10 min run time on a desktop workstation” for a MEarth star with 1000 data points (Berta et al. 2012). BLS run times of the PyKE keplbls routine have been reported as “26 min (...) using a 3 GHz Intel Core 2 Duo” (Kinemuchi et al. 2012)<sup>7</sup>. BLS speedup factors of 25× for K2 sized data are projected using optimal period sampling and optimal phase sampling (Ofir 2014) for nominal BLS sensitivity.

Longer data sets such as the long cadence (LC) light curves from the *Kepler* primary mission (K1) with ~60 000 data points and  $S = 4.25$  yr are more demanding, as are K1 short cadence (K1 SC) data. PLATO long cadence (P LC) light curves will have a 10 min sampling over 2 yr (per field of view) and about  $10^5$  cadences. PLATO short cadence (P SC) will be 25 s and deliver  $2.5 \times 10^6$  cadences over 2 yr per light curve.

All quoted TLS run times include the calculation of initial star-specific templates, which requires ~10 ms for the quadratic limb-darkening law. We explain these technicalities in more detail in the Appendix A.

### 3.5. Comparison to other transit detection algorithms

Mislis et al. (2016) claim that their machine-learning method is 1000 times faster than BLS, but do not include (or state) the run time required for the training part. The method is described to detect 8% more planets compared to BLS. However, no false/true positive/negative rates are given (as in our Fig. 6), preventing a comparison to TLS.

Training times for algorithms based on deep-learning to detect transits are of order several thousand CPU hours (Shallue & Vanderburg 2018). It remains unclear whether the training could be re-used in later searches. What is more, the authors describe a drop in model performance toward lower S/N transits, because few (real world) training candidates with low S/N were available.

Other studies find that random forest classifiers and convolutional neural networks produce a significant fraction of false-positives (Schanche et al. 2019). Depending on the threshold of a detection, it may also result in the outcome that BLS has the smallest fraction of false negatives (missed detections), 5% versus 11–14% for various machine classifiers (Pearson et al. 2018).

To fairly assess machine learning algorithms for transit detection, we recommend to perform an independent benchmark which includes BLS and TLS. This is beyond the scope of this paper but can be a natural follow-up work to it.

## 4. Discussion

### 4.1. Arbitrary signal shapes

TLS can be used with arbitrary search functions to detect other kinds of periodic events in stellar (or other) light curves. We plan to implement a user-friendly interface for such functions in the next release. As an example, although stellar flares are not known to be strictly periodic, flares from TRAPPIST-1 appear to be semi-periodic (Morris et al. 2018). An analytic description of a stellar flare (Davenport et al. 2014) could be used to search for periodically flaring stars. What is more, many phenomena related to exoplanets were not expected or known before they were found.

It would also be possible to feed TLS with an analytic description of exocometary transits (Rappaport et al. 2018;

Kennedy et al. 2019) or disintegrating planets with comet-like tails (Bochinski et al. 2015; Garai 2018), atmospheric refraction (Dalba 2017, 2018), exoplanetary rings (Barnes & Fortney 2004; Ohta et al. 2009; Tusnski & Valio 2011; Aizawa et al. 2017; Hatchett et al. 2018) or artificial shapes such as rectangles (Arnold 2005) as well as starshades at the Lagrange points (Gaidos 2017; Moores & Welch 2018).

### 4.2. Issues with uneven sampling and data gaps

TLS moves the model transit curve over the data points in phase space, very much akin to a moving window. This procedure assumes constant cadences or steps in phase space. Constant steps in phase space are only present on average, however. Transit timing variations can sometimes induce a stroboscopic effect: Cadences that are constant in time may not be constant in phase due to resonances with the observational cadence (Szabó et al. 2013). This should only affect a small fraction of all planets.

Variable cadences result in a morphologically distorted transit shapes, incorrect transit duration estimates, and usually reduce the SDE. Small variations of the cadence are negligible, for example from barycentering the timestamps of the *Kepler* satellite, which accounts for a variation of  $\lesssim 8$  min and is therefore substantially smaller than even the shortest known transit duration of ~40 min (Rappaport et al. 2013). Partially observed transits, which can occur near data gaps or near the beginning and end of observations, may cause similar issues. When the number of observed transits is large, for example more than a dozen, the effect of partial transits is small and this is (and will be) valid for the vast majority of planets detected by *Kepler* (and TESS and PLATO) due to the missions’ long duty cycles. Even in the case of contamination with partial transits, the phase-folded transit light curve is usually better approximated by our TLS default transit template than by a box.

### 4.3. Observational biases due to the transit shape template

Observational biases for transiting planets (Kipping & Sandford 2016) are partly due to the box shaped transit fit (when using BLS). Even when fitting a better transit shape like the TLS template, similar biases can be expected since our template curve cannot be a perfect fit for all transits. For example, an eccentric or V-form grazing transit shape is substantially different from a box. This causes increased noise, resulting in lower detection efficiency. A few real, but rare, transit shapes might be closer to a box than to the reference transit template, resulting in a different set of observational biases. Characterizing these can be a natural follow-up work.

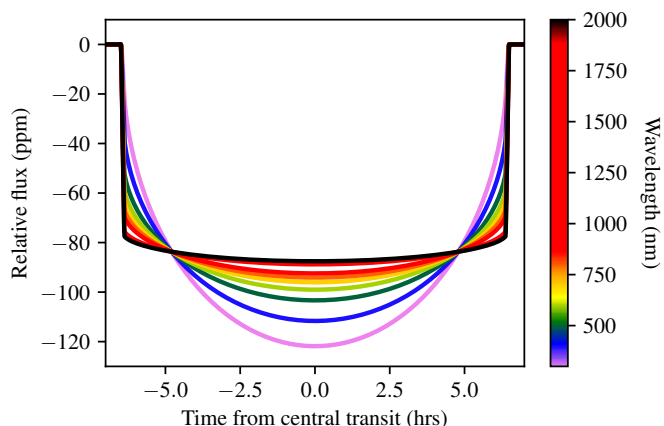
### 4.4. Correlated noise

Our definition of the SDE, extending from Eq. (1) to Eq. (5), assumes that the noise in the light curve is uncorrelated in time. For cases of correlated noise, which is often caused by stellar activity and instrumental systematics, SDE is an overestimate. The transit evaluation metrics provided by TLS (Appendix C) include the `snr_pink` and the `snr_pink_per_transit`, which can be used as an indication for correlated noise, when compared to `snr` and the `snr_per_transit`.

### 4.5. Wavelength dependence of transit shape

The advantage of TLS over BLS is largest at short wavelengths, that is in the blue-optical regime of the electromagnetic

<sup>7</sup> <https://keplergo.arc.nasa.gov/ContributedSoftwareKeplbls.shtml>



**Fig. 10.** Transit light curves for different wavelengths from the UV (300 nm) to the NIR ( $2\mu\text{m}$ ) are increasingly box-shaped. Simulation for a central Earth/sun transit using measurements from Hestroffer & Magnan (1998).

spectrum, where stellar limb darkening is most pronounced. At longer wavelengths, transits become increasingly box-shaped and the advantage of TLS over BLS vanishes. As can be seen in Fig. 10, a central Earth/sun transit observed at  $\lambda = 2\mu\text{m}$  is well approximated by a box. While TLS can be parameterized with appropriate limb-darkening for such a bandpass and thus transit shape, its advantage over BLS is reduced to the trapezoidal ingress and egress shape (and the lack of binning). The detection efficiencies of BLS and TLS essentially match at long ( $\sim 2\mu\text{m}$ ) wavelengths.

#### 4.6. Eccentricity and transit timing/duration variations

Throughout this paper, we have assumed planets in circular orbits with linear ephemerides. While this is a valid approximation for many real planets, a certain fraction of planets exhibits deviations in the form of eccentricity and/or transit timing/duration variations. All of these effects cause the signal to deviate from the TLS default template and reduce the SDE. This is also true for BLS, as these effects broaden the signal in phase space. For the vast majority of these cases, the transit shapes are still closer to the TLS default template than to a box, so that the advantage of TLS over BLS holds. Usually, TTVs are very small, less than a few percent of the transit duration. Studying other cases, where a different signal template is worthwhile, could be a natural follow-up work. A search for less common transit shapes can readily be made through TLS' interface to *batman* (Kreidberg 2015a,b), using of the underlying Mandel & Agol (2002) analytic transit model.

#### 4.7. Data binning

While the time required for phase-folding and sorting is identical for TLS and BLS, the latter is faster by an order of magnitude due to the binning of the phase-folded light curve. A box-shaped function allows for binning with minimal loss in quality and a fixed factor speed gain (Ofir 2014). For typical long cadence data (for example 30 min), however, binning cannot be recommended for TLS because it would smear the ingress and egress shapes of most transits and therefore reduce the sensitivity – if the ingress and egress duration is short. For high impact parameter transits, this duration can be long, and binning may be acceptable. In general, binning is adequate if there are many data points between

two phase grid points at the critical phase sampling. A detailed analysis of such trade-offs can be a natural part of follow-up work on TLS.

Our current TLS implementation does not compensate for morphological light-curve distortions (temporal smearing effects) due to finite integration time (Kipping 2010). It is computationally prohibitive to re-compute the transit shape template at every test period. Instead, an optimal re-computation grid of periods could be derived following Eq. (40) in Kipping (2010). This is an open feature request for TLS and can be part of natural follow-up work.

## 5. Conclusion

The default transit search function of TLS is a model transit light curve optimized to find small planets. We have constructed this template based on 2346 small ( $R_p/R_s < 0.05$ ) planets and planet candidates observed with the *Kepler* mission. With this template, or a user-specified signal shape, TLS analyses the entire, unbinned data of the phase-folded light curve. Our transit injection-retrieval experiments with white noise light curves of an Earth-sized planet around a sun-like star demonstrate that these improvements yield a 17 percentage points higher true positive rate for TLS ( $\sim 93\%$ ) compared to BLS ( $\sim 76\%$ ) if the false alarm rates are chosen to be 1%, respectively. At the same time, the TLS false negative rate (7%) is significantly smaller than that of BLS (24%). In other words, TLS is substantially more efficient and reliable in finding small planets than BLS.

The test statistic of TLS is a modified version of the signal detection efficiency (SDE) used by the standard transit detection algorithm BLS. The SDE for TLS is derived from all the data points in the phase-folded light curves and not from the binned phase-folded light curve, as done by BLS. The TLS approach is computationally more demanding but key to the increased transit detection efficiency of TLS over BLS for small planets. We also filter the SDE periodogram for a systematic noise component by dividing it through a walking median. The resulting SDE distribution for TLS yields significantly more robust detections of transit-like signals compared to BLS.

Finally, as a demonstration example for the detection of Earth-sized planets around a low-mass star, we have tested TLS with its default transit template on the K2 data of the TRAPPIST-1 system and retrieved six of seven planets together with their detection statistics and phase-folded light curves. The high detection efficiency of TLS and its optimization for computational speed makes it a natural search algorithm for small transiting planets in light curves from *Kepler*, K2, TESS, and PLATO.

*Acknowledgements.* The authors thank an anonymous referee for her or his very helpful report. RH receives funding from the German Space Agency (Deutsches Zentrum für Luft- und Raumfahrt) under PLATO Data Center grant 50001501. This research has made use of the Exoplanet Orbit Database, of the Exoplanet Data Explorer at <http://exoplanets.org>, and of NASA's ADS Bibliographic Services. The authors made use of the following software: Astropy (Astropy Collaboration 2018), NumPy (Van Der Walt et al. 2011), SciPy (Jones et al. 2001), and Matplotlib (Hunter 2007).

## References

- Aigrain, S., & Favata, F. 2002, *A&A*, 395, 625
- Aigrain, S., & Irwin, M. 2004, *MNRAS*, 350, 331
- Aigrain, S., Hodgkin, S. T., Irwin, M. J., Lewis, J. R., & Roberts, S. J. 2015, *MNRAS*, 447, 2880
- Aigrain, S., Parviainen, H., & Pope, B. J. S. 2016, *MNRAS*, 459, 2408

- Aizawa, M., Uehara, S., Masuda, K., Kawahara, H., & Suto, Y. 2017, *AJ*, **153**, 193
- Alcock, C., Allsman, R., Alves, D. R., et al. 2000, *ApJ*, **542**, 257
- Armstrong, D. J., Günther, M. N., McCormac, J., et al. 2018, *MNRAS*, **478**, 4225
- Arnold, L. F. A. 2005, *ApJ*, **627**, 534
- Astropy Collaboration (Robitaille, T. P., et al.) 2013, *A&A*, **558**, A33
- Astropy Collaboration (Price-Whelan, A. M., et al.) 2018, *AJ*, **156**, 123
- Auvergne, M., Bodin, P., Boisdard, L., et al. 2009, *A&A*, **506**, 411
- Bakos, G., Noyes, R. W., Kovács, G., et al. 2004, *PASP*, **116**, 266
- Barclay, T., Rowe, J. F., Lissauer, J. J., et al. 2013, *Nature*, **494**, 452
- Barnes, J. W., & Fortney, J. J. 2004, *ApJ*, **616**, 1193
- Berta, Z. K., Irwin, J., Charbonneau, D., Burke, C. J., & Falco, E. E. 2012, *AJ*, **144**, 145
- Bochinski, J. J., Haswell, C. A., Marsh, T. R., Dhillon, V. S., & Littlefair, S. P. 2015, *ApJ*, **800**, L21
- Bordé, P., Fressin, F., Ollivier, M., Léger, A., & Rouan, D. 2007, *Transiting Extrapolar Planets Workshop*, **366**, 145
- Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Science*, **327**, 977
- Bouffleur, R. C., Emilio, M., Pacheco, E. J., de La Reza, J. R., & da Rocha, J. C. 2014, *Formation, Detection, and Characterization of Extrasolar Habitable Planets*, 293, 410
- Brown, T. M., Latham, D. W., Everett, M. E., & Esquerdo, G. A. 2011, *AJ*, **142**, 112
- Cabrera, J., Csizmadia, S., Erikson, A., Rauer, H., & Kirste, S. 2012, *A&A*, **548**, A44
- Carter, J. A., & Agol, E. 2013, *ApJ*, **765**, 132
- Charbonneau, D., Brown, T. M., Latham, D. W., & Mayor, M. 2000, *ApJ*, **529**, L45
- Claret, A., Hauschildt, P. H., & Witte, S. 2012, *A&A*, **546**, A14
- Claret, A., Hauschildt, P. H., & Witte, S. 2013, *A&A*, **552**, A16
- Collier Cameron, A., Pollacco, D., Street, R. A., et al. 2006, *MNRAS*, **373**, 799
- Dalba, P. A. 2017, *ApJ*, **848**, 91
- Dalba, P. 2018, *American Astronomical Society Meeting Abstracts*, **231**, 128.02
- Davenport, J. R. A., Hawley, S. L., Hebb, L., et al. 2014, *ApJ*, **797**, 122
- Defaÿ, C., Deleuil, M., & Barge, P. 2001, *A&A*, **365**, 330
- Doyle, L. R., Deeg, H. J., Kozhevnikov, V. P., et al. 2000, *ApJ*, **535**, 338
- Dressing, C. D., & Charbonneau, D. 2015, *ApJ*, **807**, 45
- Ducrot, E., Sestovic, M., Morris, B. M., et al. 2018, *AJ*, **156**, 218
- Enoch, B., Haswell, C. A., Norton, A. J., et al. 2012, *A&A*, **548**, A48
- Gaidos, E. 2017, *MNRAS*, **469**, 4455
- Garai, Z. 2018, *A&A*, **611**, A63
- Gautier, III, T. N., Charbonneau, D., Rowe, J. F., et al. 2012, *ApJ*, **749**, 15
- Giacobbe, P., Damasso, M., Sozzetti, A., et al. 2012, *MNRAS*, **424**, 3101
- Gilliland, R. L., Chaplin, W. J., Dunham, E. W., et al. 2011, *ApJS*, **197**, 6
- Gilliland, R. L., Chaplin, W. J., Jenkins, J. M., Ramsey, L. W., & Smith, J. C. 2015, *AJ*, **150**, 133
- Gillon, M., Jehin, E., Lederer, S. M., et al. 2016, *Nature*, **533**, 221
- Gillon, M., Triaud, A. H. M. J., Demory, B.-O., et al. 2017, *Nature*, **542**, 456
- Hartman, J. D., & Bakos, G. Á. 2016, *Astron. Comput.*, **17**, 1
- Hatchett, W. T., Barnes, J. W., Ahlers, J. P., MacKenzie, S. M., & Hedman, M. M. 2018, *New A*, **60**, 88
- Heller, R. 2019, *A&A*, in press, DOI: 10.1051/0004-6361/201834620
- Hestroffer, D., & Magnan, C. 1998, *A&A*, **333**, 338
- Hoare, C. A. R. 1962, *Comput. J.*, **5**, 10
- Howell, S. B., Sobeck, C., Haas, M., et al. 2014, *PASP*, **126**, 398
- Huber, D., Bryson, S. T., Haas, M. R., et al. 2016, *ApJS*, **224**, 2
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- Jenkins, J. M., Doyle, L. R., & Cullers, D. K. 1996, *Icarus*, **119**, 244
- Jenkins, J. M., Caldwell, D. A., Chandrasekaran, H., et al. 2010, *ApJ*, **713**, L87
- Jenkins, J. M., Twicken, J. D., McCauliff, S., et al. 2016, in *Software and Cyberinfrastructure for Astronomy IV*, Proc. SPIE, 9913, 99133E
- Johnson, M. C., Gandolfi, D., Fridlund, M., et al. 2016, *AJ*, **151**, 171
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, *SciPy: Open Source Scientific Tools for Python*, <http://www.scipy.org/>
- Kennedy, G. M., Hope, G., Hodgkin, S. T., & Wyatt, M. C. 2019, *MNRAS*, **482**, 5587
- Kinemuchi, K., Barclay, T., Fanelli, M., et al. 2012, *PASP*, **124**, 963
- Kipping, D. M. 2010, *MNRAS*, **408**, 1758
- Kipping, D. M., & Sandford, E. 2016, *MNRAS*, **463**, 1323
- Kipping, D. M., Hartman, J., Buchhave, L. A., et al. 2013, *ApJ*, **770**, 101
- Knuth, D. 1998, *The Art of Computer Programming: Sorting and Searching*, The Art of Computer Programming (Addison-Wesley)
- Kovács, G., Zucker, S., & Mazeh, T. 2002, *A&A*, **391**, 369
- Kovács, G., Zucker, S., & Mazeh, T. 2016, *Astrophysics Source Code Library* [record ascl:1607.008]
- Kreidberg, L. 2015a, *PASP*, **127**, 1161
- Kreidberg, L. 2015b, *Astrophysics Source Code Library* [record ascl:1510.002]
- Lam, S. K., Pitrou, A., & Seibert, S. 2015, *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC – LLVM 15* (ACM Press)
- Livingston, J. H., Endl, M., Dai, F., et al. 2018, *AJ*, **156**, 78
- Luger, R., Agol, E., Kruse, E., et al. 2016, *AJ*, **152**, 100
- Luger, R., Sestovic, M., Kruse, E., et al. 2017, *Nat. Astron.*, **1**, 0129
- Mandel, K., & Agol, E. 2002, *ApJ*, **580**, L171
- Mislis, D., Bachelet, E., Alsubai, K. A., Bramich, D. M., & Parley, N. 2016, *MNRAS*, **455**, 626
- Moore, G. E. 1965, *Electronics*, **38**
- Moore, J. E., & Welch, D. 2018, *Res. Notes Am. Astron. Soc.*, **2**, 34
- Morris, B. M., Agol, E., Davenport, J. R. A., & Hawley, S. L. 2018, *ApJ*, **857**, 39
- Moutou, C., Pont, F., Barge, P., et al. 2005, *A&A*, **437**, 355
- Ofir, A. 2008, *MNRAS*, **387**, 1597
- Ofir, A. 2014, *A&A*, **561**, A138
- Ohta, Y., Taruya, A., & Suto, Y. 2009, *ApJ*, **690**, 1
- Osborn, H. P., Santerne, A., Barros, S. C. C., et al. 2017, *A&A*, **604**, A19
- Pearson, K. A., Palafox, L., & Griffith, C. A. 2018, *MNRAS*, **474**, 478
- Pepper, J., Pogge, R. W., DePoy, D. L., et al. 2007, *PASP*, **119**, 923
- Petigura, E. A., Marcy, G. W., & Howard, A. W. 2013, *ApJ*, **770**, 69
- Pollacco, D. L., Skillen, I., Collier Cameron, A., et al. 2006, *PASP*, **118**, 1407
- Pont, F., Zucker, S., & Queloz, D. 2006, *MNRAS*, **373**, 231
- Pope, B. J. S., Parviainen, H., & Aigrain, S. 2016, *MNRAS*, **461**, 3399
- Rappaport, S., Sanchis-Ojeda, R., Rogers, L. A., Levine, A., & Winn, J. N. 2013, *ApJ*, **773**, L15
- Rappaport, S., Vanderburg, A., Jacobs, T., et al. 2018, *MNRAS*, **474**, 1453
- Rauer, H., Catala, C., Aerts, C., et al. 2014, *Exp. Astron.*, **38**, 249
- Régulo, C., Almenara, J. M., Alonso, R., Deeg, H., & Roca Cortés, T. 2007, *A&A*, **467**, 1345
- Régulo, C., Almenara, J. M., & Deeg, H. J. 2009, *Transiting Planets*, **253**, 374
- Renner, S., Rauer, H., Erikson, A., et al. 2008, *A&A*, **492**, 617
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2014, *Space Telescopes and Instrumentation 2014: Optical, Infrared, and Millimeter Wave*, 9143, 914320
- Rodenbeck, K., Heller, R., Hippke, M., & Gizon, L. 2018, *A&A*, **617**, A49
- Sanchis-Ojeda, R., Rappaport, S., Winn, J. N., et al. 2014, *ApJ*, **787**, 47
- Savitzky, A., & Golay, M. J. E. 1964, *Anal. Chem.*, **36**, 1627
- Schanche, N., Collier Cameron, A., Hébrard, G., et al. 2019, *MNRAS*, **483**, 5534
- Schwarzenberg-Czerny, A., & Beaulieu, J. P. 2006, *MNRAS*, **365**, 165
- Seager, S., & Mallén-Ornelas, G. 2003, *ApJ*, **585**, 1038
- Shallue, C. J., & Vanderburg, A. 2018, *AJ*, **155**, 94
- Siverd, R. J., Beatty, T. G., Pepper, J., et al. 2012, *ApJ*, **761**, 123
- Stellingwerf, R. F. 1978, *ApJ*, **224**, 953
- Szabó, R., Szabó, G. M., Dálya, G., et al. 2013, *A&A*, **553**, A17
- Tal-Or, L., Mazeh, T., Alonso, R., et al. 2013, *A&A*, **553**, A30
- Tingley, B. 2003a, *A&A*, **403**, 329
- Tingley, B. 2003b, *A&A*, **408**, L5
- Tusnski, L. R. M., & Valio, A. 2011, *ApJ*, **743**, 97
- Vanderburg, A., Latham, D. W., Buchhave, L. A., et al. 2016, *ApJS*, **222**, 14
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *Comput. Sci. Eng.*, **13**, 22
- Van Grootel, V., Fernandes, C. S., Gillon, M., et al. 2018, *ApJ*, **853**, 30
- Wells, R., Poppenhaeger, K., & Watson, C. A. 2018, *MNRAS*, **473**, L131
- Wright, J. T., Fakhouri, O., Marcy, G. W., et al. 2011, *PASP*, **123**, 412
- Zhang, H., Yu, Z., Liang, E., et al. 2019, *ApJS*, **240**, 17
- Zucker, S., & Giryas, R. 2018, *AJ*, **155**, 147

## Appendix A: Optimization for computational speed

**Table A.1.** TLS run times (in minutes) for *Kepler* K1 LC data,  $R = R_{\odot}$ ,  $M = M_{\odot}$ ,  $\Delta t_0 = 0$ .

$\delta_{\text{cut}}$ (ppm)	$\pm 0.1\rho$	$\pm 0.2\rho$	No prior
10	30	50	400
50	20	40	300
100	7	10	30

**Notes.** For reference: *Astropy* BLS run time: 13 min.

A straightforward implementation of the mathematical framework described in Sect. 2 into an algorithm would result in huge computational demands with days of run time per K2 light curve for a reasonably dense grid of  $P, t_0, \delta, d$  trial values. We thus implement a computationally optimized version that produces identical results to a straightforward coding. Most aspects of our optimization are time-memory trade-offs, where repetitive calculations are identified and stored in memory after their first calculation. A memory read is often faster than a repeated calculation. The memory size requirements for the computer’s random access memory (RAM) are of order 50 MB per thread, dominated by the buffered signal shapes.

Most of the data points in a light curve are out of transit. As a consequence, the out-of-transit data account for the majority of computation time. Most important, the out-of-transit data of both the observed and of the modeled light curves are identical for a fixed trial epoch, transit duration, and orbital period. Hence, we design TLS to calculate the squared residual values of the out-of-transit data only once in the  $(P-t_0-d)$  3D parameter hyperspace of our 4D search grid. Consequently, only a small amount of in-transit squared residuals need to be calculated for the various trial depths of the transits.

We also found that instead of calculating the oversampled transit model for each  $\chi^2$  test in the  $d$  space, it is significantly faster to pre-compute all oversampled transit models for all trial durations but only for single, arbitrary transit depth for any given  $d$  value. Re-scaling the transit in  $\delta$  only requires one multiplication per data point. Resampling a model in width, however, would be considerably more expensive due to the necessary oversampling. Thus, the transit models are pre-computed and cached for any given transit duration and then they are scaled in depth on-the-fly as TLS searches through the transit duration grid.

Phase-folding involves sorting the orbital phases  $\phi_i = t_i/P$ , where  $t_i$  are the times of observations. Sorting can be extremely demanding computationally. A typical K2 light curve worth 4000 data points takes 0.5 ms to fold and sort on an Intel Core i7-7700K. The required time for folding and sorting 20 000 trial periods and 4000 trial phases would be  $\sim 11$  h. Hence, as in several BLS implementations, TLS only folds and sorts the phases once per trial period, which takes  $\sim 5$  s and implies a speed gain of a factor of 4000. We find that the fastest algorithm to sort phase arrays is “Mergesort” (von Neumann 1945, unpublished), which is typically  $\sim 20\%$  faster than the commonly used Quicksort algorithm (Hoare 1962).

The TLS reference implementation is written in pure python code, which is an interpreter-based programming language and thus comes with a speed loss. We therefore chose to implement many of the time-critical parts of TLS with the specialized numba package (Lam et al. 2015) that translates the python code in

**Table A.2.** TLS run times (in minutes) as before, but  $\delta_{\text{cut}} = 10$  ppm.

$\Delta t_0$	$\pm 0.1\rho$	$\pm 0.2\rho$	No prior
0	30	50	400
0.01	7	9	40
0.1	7	7	7

**Notes.** For reference: *Astropy* BLS run time: 13 min.  $\Delta t_0$  is the step in phase space in units of transit duration.

machine code. This procedure is called “just-in-time” compilation and saves us two orders of magnitude in computing time.

TLS can be adjusted for adequate performance on large data sets with minor compromises. When speed is critical, we recommend to first perform a search using fast binned BLS, which is adequate to recover all high-S/N planets. Using iterative runs, these significant signals can be found and masked from the data, reducing overall variance and data volume. When no significant BLS signal remains, the search can be switched to TLS.

The computational speed of TLS can then be increased with a sensible threshold for the shallowest transit signal that is fitted. For example, instrumental and stellar noise limit the detectability of signals by *Kepler* to transit depths of  $\sim 100$  ppm for periods of  $\sim 365$  d around a G2V-star, roughly the Earth-equivalent that the *Kepler* mission was originally designed for. Allowing the algorithm to fit signals of down to 10 ppm will then only result in a very high computational load, but almost certainly not in the discovery of a real 10 ppm transit signal. As the S/N of a planet is a complex combination of many factors, we have set the default TLS parametrization to a threshold of 10 ppm. For reference, the shallowest known transit is 11.9 ppm (Kepler-37 b, Barclay et al. 2013), a discovery that was made possible due to the short orbital period (13.4 d) and the stacking of  $\sim 100$  transits. For long-period planets, one may choose thresholds of, for example, 50 ppm for *Kepler* data or 100 ppm for K2 data, which avoids fitting out the complete noise floor. For reference, an Earth-sized planet transiting a dG2 star has a transit depth of  $\sim 100$  ppm. Then TLS run times are similar to those of BLS. As a follow-up work, one might develop a heuristic of a shallowest transit depth to be fitted, as a function of noise in the data, period, and other factors.

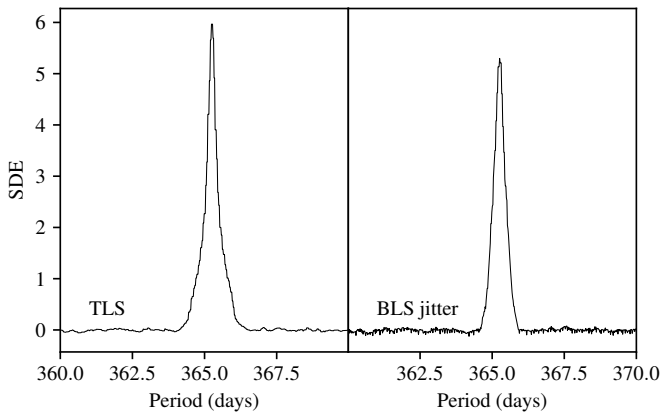
For *Kepler* K1 LC data (30 min cadence over 4.25 yr), we have tested TLS run times for different combinations of stellar density priors (from stellar radii and masses), shallowest transit fit depths ( $\delta_{\text{cut}}$ ), and phase space sampling  $\Delta t_0$  (Tables A.1, and A.2). The KIC and EPIC catalogs (Brown et al. 2011; Huber et al. 2016) typically have relative mass and radius uncertainties of 5%, so that a range of  $\pm 0.1\rho$  ( $\pm 0.2\rho$ ) gives a  $2\sigma$  ( $4\sigma$ ) confidence interval. Priors decrease run times by an order of magnitude in case of complete phase space sampling. Then, the influence of  $\delta_{\text{cut}}$  is a factor of a few between 100 ppm and 10 ppm for typical K1 data, where the standard deviation per data point (after detrending) is typically in the hundreds of ppm.

TLS also offers the option of not sampling the phase space at every cadence. For example, K1 LC data (60 000 points) allows for transit signal lengths of up to 7200 points for  $T_{14}/P = 0.12$ . Considering noise levels in the hundreds of ppm per point, shifting this long template point by point is pointless. Instead, TLS can shift the data in phase space by a user-defined fraction of the transit duration, the latter of which is measured in cadences. As an example, the default value of 1% shifts a transit signal of 200 points length by 2 points in each trial, saving 50% of the computational time. This procedure results in much faster run times because most of the computational effort is spent to

test very long transit duration. Empirically, we find that setting  $\Delta t_0 = 0.01$  (instead of zero) allows for virtually identical SDE values in *Kepler* K1 and K2 data, while  $\Delta t_0 = 0.1$  results in a detection efficiency loss of a few percent and should be used only for high S/N transits, for example, using an iterative search.

Our TLS implementation leverages all available CPU cores and shows continuous updates of the estimated remaining time and a progress bar. The user can use the estimate to balance run time and search depth.

## Appendix B: Edge effect jitter in BLS which leads to additional noise



**Fig. B.1.** Edge effect jitter in BLS (*right*) and absence of this jitter in TLS (*left*).

The original BLS implementation did not account for transit events occurring to be divided between the first and the last bin of the folded light curve. This was noted by Peter R. McCullough in 2002, and an updated version of BLS was made (`ee-bls.f`) to account for this edge effect. The patch is commonly realized by extending the phase array through appending the first bin once again at the end, so that a split transit is stitched together, and present once in full length. The disadvantage of this approach has apparently been ignored: The test statistic is affected by a small amount of additional noise. Depending on the trial period, a transit signal (if present) is sometimes partly located in the first

and the second bin. The lower (in-transit) flux values from the first bin are appended at the end of the data, resulting in a change of the ratio between out-of-transit and in-transit flux. There are phase-folded periods with one, two, or more than two bins which contain the in-transit flux. This causes a variation (over periods) of the summed noise floor, resulting in additional jitter in the test statistic. For typical *Kepler* light curves, the reduction in detection efficiency is comparable to a reduction in transit depth of  $\sim 0.1\text{--}1\%$ . TLS corrects this effect by subtracting the difference of the summed residuals between the patched and the non-patched phased data. A visualization of this effect on the statistic is shown in Fig. B.1, using synthetic data. In real data, the effect is usually overpowered by noise, and was thus ignored, but it is nonetheless present.

## Appendix C: Transit evaluation metrics

In addition to the SR and SDE transit search statistic (Sect. 2.2), our python implementation of TLS outputs the period of the highest SDE value, its corresponding  $t_0$ ,  $\delta$ ,  $d$ , and the planet-to-star radius ratio for zero transit impact parameter  $R_P/R_\star = \sqrt{\delta(1 - c_1/3 - c_2/6)}$  (Heller 2019), where  $c_1$  and  $c_2$  are the limb darkening coefficients of a quadratic limb darkening law. TLS also offers a range of automated evaluation parameters of the detected transits such as

- the ratio of the signal to the white noise of the stacked transits (`snr`);
- the ratio of the signal to the pink noise of the stacked transits (`snr_pink`) (Pont et al. 2006; Hartman & Bakos 2016);
- the ratio of the signal to the white noise of the individual transits (`snr_per_transit`);
- the ratio of the signal to the white noise of the individual transits (`snr_pink_per_transit`);
- the significance (in units of standard deviations) between the depths of the odd and even transits (`odd_even_mismatch`);
- the number of transits with in-transit data points (`distinct_transit_count`);
- the number of transits with no in-transit data points (`empty_transit_count`);
- the total number of transits (`transit_count`);
- the number of data points for each transit (`per_transit_count`).

Our online release of TLS comes with a documentation of these parameters and of additional evaluation metrics.