A&A, 686, A239 (2024)
https://doi.org/10.1051/0004-6361/202348951
© The Authors 2024

Astronomy & Astrophysics

Detailed cool star flare morphology with CHEOPS and TESS⋆,⋆⋆,⋆⋆⋆

G. Bruno[1], I. Pagano[1], G. Scandariato[1], H.-G. Florén[2], A. Brandeker[2], G. Olofsson[2], P. F. L. Maxted[3], A. Fortier[4,5], S. G. Sousa[6], S. Sulis[7], V. Van Grootel[8], Z. Garai[9,10,11], A. Boldog[12,13,9], L. Kriskovics[12,13], Gy. M. Szabó[10,9], D. Gandolfi[14], Y. Alibert[5,4], R. Alonso[15,16], T. Bárczy[17], D. Barrado Navascues[18], S. C. C. Barros[6,19], W. Baumjohann[20], M. Beck[21], T. Beck[4], W. Benz[4,5], N. Billot[21], L. Borsato[22], C. Broeg[4,5], A. Collier Cameron[23], Sz. Csizmadia[24], P. E. Cubillos[25,20], M. B. Davies[26], M. Deleuil[7], A. Deline[21], L. Delrez[27,8,28], O. D. S. Demangeon[6,19], B.-O. Demory[5,4], D. Ehrenreich[21,29], A. Erikson[24], J. Farinato[22], L. Fossati[20], M. Fridlund[30,31], M. Gillon[27], M. Güdel[32], M. N. Günther[33], A. Heitzmann[21], Ch. Helling[20,34], S. Hoyer[7], K. G. Isaak[33], L. L. Kiss[35,36], K. W. F. Lam[24], J. Laskar[37], A. Lecavelier des Etangs[38], M. Lendl[21], D. Magrin[22], C. Mordasini[4,5], V. Nascimbeni[22], R. Ottensamer[32], E. Pallé[15,16], G. Peter[39], G. Piotto[22,40], D. Pollacco[41], D. Queloz[42,43], R. Ragazzoni[22,40], N. Rando[33], F. Ratti[33], H. Rauer[24,44], I. Ribas[45,46], N. C. Santos[6,19], M. Sarajlic[3], D. Ségransan[21], A. E. Simon[4,5], V. Singh[1], A. M. S. Smith[24], M. Stalport[8,27], N. Thomas[4], S. Udry[21], B. Ulmer[39], J. Venturini[21], E. Villaver[15,16], N. A. Walton[47], and T. G. Wilson[41]

(Affiliations can be found after the references)

Received 13 December 2023 / Accepted 11 March 2024

ABSTRACT

Context. White-light stellar flares are proxies for some of the most energetic types of flares, but their triggering mechanism is still poorly understood. As they are associated with strong X and ultraviolet emission, their study is particularly relevant to estimate the amount of high-energy irradiation onto the atmospheres of exoplanets, especially those in their stars' habitable zone.

Aims. We used the high-cadence, high-photometric capabilities of the CHEOPS and TESS space telescopes to study the detailed morphology of white-light flares occurring in a sample of 130 late-K and M stars, and compared our findings with results obtained at a lower cadence.

Methods. We employed dedicated software for the reduction of 3 s cadence CHEOPS data, and adopted the 20 s cadence TESS data reduced by their official processing pipeline. We developed an algorithm to separate multi-peak flare profiles into their components, in order to contrast them to those of single-peak, classical flares. We also exploited this tool to estimate amplitudes and periodicities in a small sample of quasi-periodic pulsation (QPP) candidates.

Results. Complex flares represent a significant percentage (≳30%) of the detected outburst events. Our findings suggest that high-impulse flares are more frequent than suspected from lower-cadence data, so that the most impactful flux levels that hit close-in exoplanets might be more time-limited than expected. We found significant differences in the duration distributions of single and complex flare components, but not in their peak luminosity. A statistical analysis of the flare parameter distributions provides marginal support for their description with a log-normal instead of a power-law function, leaving the door open to several flare formation scenarios. We tentatively confirmed previous results about QPPs in high-cadence photometry, report the possible detection of a pre-flare dip, and did not find hints of photometric variability due to an undetected flare background.

Conclusions. The high-cadence study of stellar hosts might be crucial to evaluate the impact of their flares on close-in exoplanets, as their impulsive phase emission might otherwise be incorrectly estimated. Future telescopes such as PLATO and Ariel, thanks to their high-cadence capability, will help in this respect. As the details of flare profiles and of the shape of their parameter distributions are made more accessible by continuing to increase the instrument precision and time resolution, the models used to interpret them and their role in star-planet interactions might need to be updated constantly.

Key words. methods: data analysis – techniques: photometric – stars: activity – stars: flare – planetary systems

1. Introduction

As an outcome of astrophysical dynamos, stellar activity manifests itself under a variety of timescales and wavelengths, and related observables are produced from the photosphere up until the corona. From the day- to month-, or even year-long cycles due to starspots, faculae, and their evolution (e.g. Hall 1991; Lanza et al. 1998; Namekata et al. 2019, 2020), down to impulsive events such as stellar flares and coronal mass ejections (CMEs) that are generated in active regions (e.g. Hudson 1991; Maehara et al. 2012; Walkowicz et al. 2011; Kowalski et al. 2013), these phenomena are among the main factors that shape the stellar energy balance and the environment where planets live. In this respect, the impact of flares and CMEs on planetary atmospheres, including their possible erosion, ionisation, and triggering of photochemical reactions in their upper layers,

⋆ The CHEOPS photometry discussed in this paper is available in electronic form at the CDS via anonymous ftp to cdsarc.cds.unistra.fr (130.79.128.5) or via https://cdsarc.cds.unistra.fr/viz-bin/cat/J/A+A/686/A239.

⋆⋆ This study used CHEOPS data observed as part of the Guaranteed Time Observation programmes CH_PR100018 (PI: I. Pagano) and CH_PR100010 (PI: G. Szabó).

⋆⋆⋆ The main parts of the code we developed can be found on GitHub (https://github.com/giovbruno/flarefinder).

Open Access article, published by EDP Sciences, under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.
This article is published in open access under the Subscribe to Open model. Subscribe to A&A to support open access publication.

has recently attracted attention from many teams (e.g. Sanz-Forcada et al. 2010; Venot et al. 2016; Spake et al. 2018; Rodríguez-Barrera et al. 2018; Guilluy et al. 2020; Chen et al. 2021; Locci et al. 2022; Colombo et al. 2022; Maggio et al. 2022, 2023; Jackman et al. 2023; Louca et al. 2023). The interest in these topics was increased by the discovery of planets in the habitable zone of M stars (e.g. Gillon et al. 2017), whose high-energy irradiation is both able to stimulate the formation of biologically valuable molecules and to destroy the chemical bonds that make life as we know it possible (e.g. Rimmer et al. 2018; Barth et al. 2021). Another key parameter is the flare rate: frequent events are likely to increase their effect on exoplanet atmospheres, in addition to (when rocky) surface and even interiors, as they do not leave the time necessary for the induced perturbations to fade out (Venot et al. 2016; Vida et al. 2017; Hazra et al. 2020; Airapetian et al. 2020; Grayver et al. 2022; Nicholls et al. 2023).

Flares on M stars are known to often reach several orders of magnitude higher energies than solar flares, which hardly go beyond 10^{32} erg (Maehara et al. 2012; Shibayama et al. 2013; Loyd et al. 2018). While flare emission extends from gamma rays to the radio, the most energetic events are those that can be observed in the optical, and happen when the concomitant extreme ultraviolet (EUV) and soft X-ray luminosity reaches particularly high levels (McIntosh & Donnelly 1972; Neidig 1983). Hence, the study of white-light flares, which behave similarly in the Sun and other stars (Neidig 1989), is key to understand the impact of stellar activity on planetary atmospheres and magnetospheres, including that of the Earth (e.g. Airapetian et al. 2016).

At the lowest energies, flares are not less interesting. Through reconstructions of the solar disc at about 1 min cadence in EUV wavelengths, it was possible to detect solar micro- and nanoflares (Shimizu & Tsuneta 1997; Aschwanden et al. 2000), defined by an emitted energy in the range 10^{24}–10^{27} erg and $\lesssim 10^{24}$ erg, respectively. On other stars, such low-energy flares still lack observational confirmation (e.g. Yang et al. 2017). Assuming a similar flare generation mechanism, we would expect flares on M dwarfs to reach equally low energies. Despite the intriguing idea that these processes, if frequent enough, might provide a solution to the coronal heating problem (Parker 1988; Haisch et al. 1991; Hudson 1991), no evidence to date definitively supports nanoflares against competing hypotheses (Parnell & De Moortel 2012; Dillon et al. 2020; Bogachev & Erkhova 2023). In this respect, disc-integrated optical photometry does not provide conclusive information, as nanoflares are supposed to occur in the quiet corona and to be hidden in the optical measurement noise (e.g. Benz 2017, and references therein). This is one of the reasons why it is still unclear whether the distributions of the nanoflare and larger energy outburst follow similar trends (e.g. Maehara et al. 2015; Aschwanden 2022b).

In single-filter, visible photometry, stellar flares have been observed down to cadences of minutes and tens of seconds. To date, the largest surveys of high-cadence white-light flaring stars have been carried out thanks to Transiting Exoplanet Survey Satellite (TESS) and the Next Generation Transit Survey (NGTS), with cadences of 20 s for the former (Howard & MacGregor 2022), and 10 s for the latter (Jackman et al. 2020, 2023). While increasing the time resolution from 30 to 1 min has not highlighted significant differences in the flare-emitted energy in Kepler data (Raetz et al. 2020), the energy output of the widely varied 20 s cadence flare profiles has only started to be explored.

The observation of M dwarf white-light flares at 20 s cadence revealed that a large fraction of outbursts has a non-classical profile (Howard & MacGregor 2022). Multi-peak flare shapes were found to be the norm, and could indicate cascades of emissions from a single active region or sympathetic flares from adjacent regions (Hawley et al. 2014; Davenport 2016; Schrijver & Higgins 2015; Kowalski et al. 2019). Such new data sets also allowed the detection of magneto-hydrodynamic (MHD) quasi-periodic pulsations (QPPs) in the stellar plasma down to a few minutes period, thereby offering new indications about their potential coronal heating role (Nakariakov & Melnikov 2009; Zimovets et al. 2021) and on their impact onto exoplanet atmospheres (e.g. Ramsay et al. 2021).

An encompassing picture of the physical avenues that trigger stellar flares requires detailed observations of the events occurring before the flare rise phase. Several authors have reported both photometric and spectroscopic reductions in the stellar flux before the impulsive phases of a few dMe stellar flares, with most of the increased absorption in the ultraviolet (UV, Rodono et al. 1979; Giampapa et al. 1982; Doyle et al. 1988; Peres et al. 1993; Ventura et al. 1995; Zalinian et al. 2002; Leitzinger et al. 2014). Such 'dips' or 'black-light flares' were found to happen within the half-hour prior to the flare impulsive phase, have amplitude from 1 up to 20% of the quiescent flux level, and duration from a few seconds up to a few tens of minutes. Attempted explanations connect these dips to details of the flare generation mechanism, and suggest these events might happen on the Sun, too (Henoux et al. 1990; Aboudarham & Henoux 1987; van Driel-Gesztelyi et al. 1994; Tovmassian et al. 2003). Stars with reported dips do not belong to the same sub-stellar type, as they span the M1 to M5 types, and are both isolated and in dMe binary systems. Additionally, dips do not anticipate particularly energetic flares: the one observed by Ventura et al. (1995), for example, emitted an energy of $\simeq 10^{32}$ erg in a no-filter optical bandpass, while superflares are characterised by energies of at least 10^{33} erg. However, the scarcity of observed events and the variety of characteristics of such dips hampered the possibility of comparing models and building a unified picture.

The Characterising ExOPlanet Satellite (CHEOPS) space telescope (Benz et al. 2021) has demonstrated exquisite photometric precision at 3 s cadence (Morgado et al. 2022), and so represents a unique opportunity to open an even newer window on stellar flares. In this study, we addressed the details of complex white-light flare morphology, looked for low-energy outbursts, searched for pre-flare dips and QPPs in a sample of 130 late-K and M dwarf stars observed as part of CHEOPS's ancillary science programme. We took advantage of the highest time cadence that the CHEOPS and TESS (Ricker et al. 2015) space telescopes can achieve, which enables a comparable precision in the detection of a few second to a few tens of second details in optical light curves, respectively. In Sect. 2, we describe our sample. In Sect. 3, we discuss the details of our data sets and the reduction techniques we adopted. Section 4 is dedicated to the modelling of the flare candidate profiles and Sect. 5 to the experiments made to test the algorithm we developed for our analysis. In Sect. 6 we present our results, which we discuss as we draw our conclusions in Sect. 7.

2. Target selection

Among the CHEOPS Guaranteed Time Observers operations, an Ancillary Science programme has been dedicated to monitoring the short-timescale micro-variability of main sequence late-K to M dwarfs. Some of our targets are stars which were considered for radial velocity exoplanet searches, for example, with the Calar Alto high-Resolution search for M dwarfs with

Exoearths with Near-infrared and optical Échelle Spectrographs (CARMENES, Alonso-Floriano et al. 2015; Cortés-Contreras et al. 2017). Others were selected from the M dwarf-related literature: for instance, stars for which pre-flare dips were detected. The list of our 130 K5V to M 5V targets, and their relevant parameters which are available in the literature, can be found in Table A.1.

We adopted the *Gaia G* band (Gaia Collaboration 2023) magnitude of the targets, and assessed their distance using the Early Data Relase 3 measurements by Bailer-Jones et al. (2021). When available, stellar effective temperatures ($T_{\rm eff}$) were obtained from the PASTEL catalogue (Soubiran et al. 2010), otherwise from a more general literature search and comparison with the spectral type reported on SIMBAD (Wenger et al. 2000).

We also searched the literature for activity indicators. We inspected the Strasbourg astronomical Data Center (CDS)[1] for published values of X and UV emission, $\log R'_{\rm HK}$, and $v \sin i$, and found this latter parameter to be the most frequently available (122 out of 130 targets). When different $v \sin i$ values were reported, we adopted their mean, and found $v \sin i < 5$ km s^{-1} for 80% of the objects. Among those targets with an available $\log R'_{\rm HK}$ (33 out of 130), we found 67% to have a value <−4.8. As we were only interested in a general indication of the activity level of our targets, we did not correct the $\log R'_{\rm HK}$ values for interstellar extinction. We did not attempt to measure stellar rotation periods from TESS data, as the 28-days duration of TESS light curves is likely too short to capture a full rotation for most of the late-type stars in our sample (McQuillan et al. 2014).

Histograms for the retrieved stellar parameters are illustrated in Fig. 1, and are consistent with a mostly low activity level for our objects. This was expected, given that most of our targets were selected from exoplanet-search stellar samples.

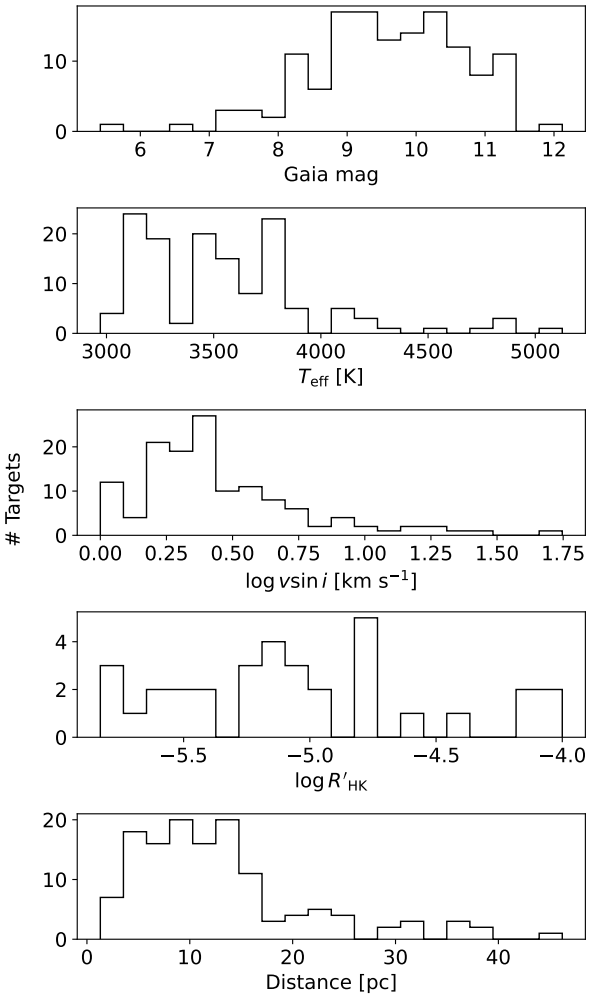

Fig. 1. Parameters for our stellar sample. From top to bottom: histograms for the *Gaia* magnitude, published $T_{\rm eff}$, $v \sin i$, $\log R'_{\rm HK}$ values and distance.

3. Observations and data reduction

The main steps of our flare search were (1) detrending of each light curve to evaluate the 'quiet' stellar flux level; (2) identification of flare peak candidates; (3) flare validation and profile fit; (4) comparison of the flare profile fit with a model including a pre-flare flux drop and inspection for QPPs in the fit residuals. The analysis was carried out with an automatic algorithm that we developed for this specific purpose. We performed a visual inspection of the results, which helped us tune the parameters of our algorithm to better suit the different types of data. In both cases, we found steps (1) and (3) to be the most critical and the most prone to errors.

3.1. CHEOPS light curves

This paper reports on the full extent of the ancillary science programme dedicated to late-K and M dwarf flares (CH_PR100018, PI: I. Pagano), active between the beginning of scientific observations in April 2020 and September 2023, totalling ~95 days spent on target. Further AU Mic observations from programme CH_100010 (PI: G. Szabó) were used, for an additional ~11 days observation time.

To study the detailed morphology of flares, photometric precision is as important as time resolution. Our goal was to obtain a flux measurement every few seconds; however, the CHEOPS standard mode of operations for $V \gtrsim 9$ stars (i.e. most of our targets) delivers images called 'subarrays' every few tens of seconds, as a result of on-board stacking of shorter exposures. These shorter exposures are not lost, but are recorded and downlinked as 50-pixel diameter 'imagettes' (Benz et al. 2021). Working with imagettes involves a partial loss of information compared to subarrays, which have a 200-pixel diameter (e.g. about each exposure's background), so that these latter are still useful for the correction of instrumental effects. The imagette exposure time can be adjusted and affects the subarray exposure time: for our operations, we found a 3 s cadence to be a good compromise between the actual on-target time and the telescope duty cycle. CHEOPS can reach as low a cadence as 1 ms, but this is only useful for $V < 6$ targets and would have implied a suboptimal combination of on-target time and duty cycle for our targets.

Imagette data sets are not automatically processed by the Data Reduction Pipeline (DRP, Hoyer et al. 2020), so that we relied on an ad hoc tool for imagette reduction: the Data Reduction Tool (DRT), described in Morgado et al. (2022) and Fortier et al. (2024). After processing and for every visit, we rejected those points with a background flux deviating by more than 5σ from its median value, and clipped the measurements with a flux value $>5\sigma$ lower than a smoothed version of the light curve. In order to avoid removing data points belonging to flares, we did not carry out any clipping above the smoothed flux level.

We normalised the light curve of each visit with a low-order polynomial as a function of time, so to remove any long-period

[1] http://cds.unistra.fr/

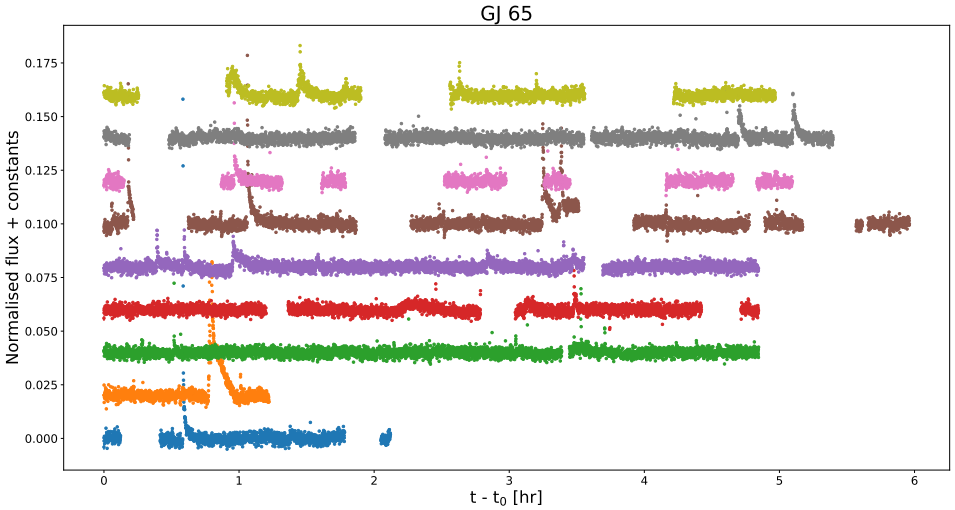

Fig. 2. GJ 65 CHEOPS DRT imagette light curves, detrended for roll-angle effects and stellar variability. Each CHEOPS visit is assigned a different colour and shifted for clarity.

trend likely due to stellar activity, with GJ 65 and AU Mic being the only targets which required a higher degree. We tested polynomials with degrees from three to ten and, for each degree, iteratively sigma-clipped the fit residuals until no more residual data points were above the clipping threshold. This allowed us to minimise the impact of the flares in the detrending procedure. After this, the polynomial fit with the lowest Akaike Information Criterion (AIC, Akaike 1974) was adopted.

Because of the irregular shape of CHEOPS's point spread function and the nadir-locked rotation of the telescope aimed at maintaining thermal stability, the satellite's observations are affected by contamination from nearby stars which are in phase with its $\varphi \in [0, 2\pi)$ roll angle (e.g. Lendl et al. 2020). To detrend the data from this instrumental effect, we combined all observations for a given target to maximise their signal-to-noise ratio (S/N) in the roll-angle space. Given the constraints imposed by the filler nature of the programme, we found a duration of three CHEOPS orbits (i.e. ∼300 min) for each observing sequence or 'visit' to be an affordable compromise. After the first months of operations, we therefore set this duration for all the visits of our programme. Also, merging all visits for a given target minimised the data gaps due to the passage of the telescope across the South Atlantic Anomaly (Benz et al. 2021), and amounting to up to 50% of each visit observation time.

We then followed Scandariato et al. (2017) to fit a model Θ for roll-angle systematics with a combination of sines and cosines, that is,

$$\Theta(\varphi) = \sum_{i=1}^{5} \left[a_i \sin(i \cdot \varphi) + b_i \cos(i \cdot \varphi) \right], \quad (1)$$

where a_i and b_i are the i-th coefficients in the Fourier reconstruction of the sinusoidal signal. After experimenting with the number of harmonics, we found i up to 5 was able to provide sufficient flattening of the residuals. The standard deviation of the flattened light curve, without outliers and flare candidates, was used as the data point uncertainty in the flattened light curve. The result of the detrending for the light curves of GJ 65, where both instrumental effects and stellar long-period variability were removed, is shown in Fig. 2.

Our detrending procedure was effective for most datasets, but less so for the few light curves that were contaminated by a Solar System body lying within $\lesssim 25°$ from the target position, or by particularly bright neighbouring objects. The resulting strong non-periodicity of the roll-angle signal in different visits required a visual inspection of the detrended light curves, the masking of some portions of the data, and the discarding of false positive detections. In particular, we rejected the light curves of 27 out of 130 targets (21%), where however we could not identify any evident flare by visual inspection.

3.2. TESS light curves

The TESS telescope uses a similar, even if redder, photometric band to CHEOPS[2] Its full-sky coverage, photometric precision, and 20 s cadence mode made available since its Cycle 3 (Sector 27 onward), was relevant for our goals, as it allowed long, continuous, and high-cadence observations of several of our targets. We therefore downloaded from the Mikulski Archive for Space Telescopes[3] all available 20 s TESS light curves for our targets at the time of writing, or 106 light curves for 73 objects, until Sector 67. The list of programmes, and the corresponding PIs, from which the data were obtained can be found in Table A.2.

We used the Pre-search Data Conditioning Simple Aperture Photometry (PDCSAP) flux, where instrumental trends are removed by the TESS pipeline (Jenkins et al. 2016). As reported in a few cases, PDCSAP correction might add spurious effects to the data (e.g. Nardiello et al. 2022); in our case, however, the trend and flare timescales are different enough to prefer this latter over the Simple Aperture Photometry (SAP) format, where instrument artefacts are not removed.

We first derived a flattened version of each light curve by removing its respective smoothed version. The smoothing function was computed with a Hann window, with a window length of ∼0.2 times the duration of the most likely periodicity in the light curve. This latter was found using ASTROPY's implementation of the Lomb-Scargle periodogram (Astropy Collaboration 2013, 2018, 2022). The smoothing was carried out iteratively, by sigma-clipping the outliers from the smoothing process at each iteration and stopping when no more outliers were present, and therefore minimising the contribution of the flares to the detrending process.

[2] A comparison of their bandpasses can be found at this page
[3] https://archive.stsci.edu/tess/bulk_downloads/bulk_downloads_ffi-tp-lc-dv.html

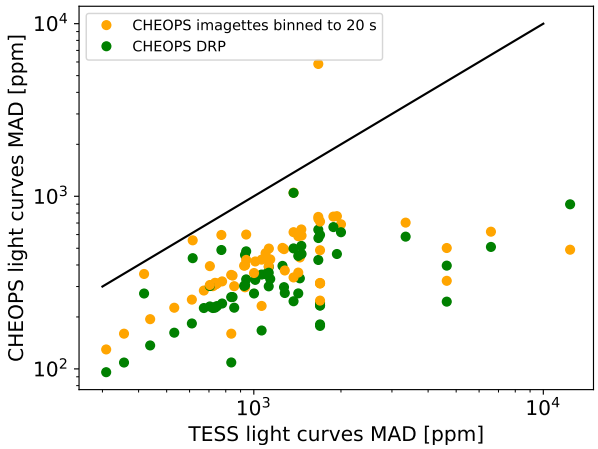

Fig. 3. Median absolute deviation (MAD) of the CHEOPS subarray (DRP) and imagette (DRT) light curves (binned to 20 s), as a function of the one of the 20 s TESS light curves. Only targets with both CHEOPS and TESS data are represented. The one-to-one line is marked in black.

We also tested other approaches to smooth the light curves, such as pre-whitening and Gaussian Process regression. In all cases, we found that the smoothing resulted in incorrectly increased smoothing function values before and after the most energetic flares, so that their flattened profiles were dampened. This also caused spurious pre- and post-flare flux drops to appear. We took this into account by determining correction factors on our results based on injection tests (Sect. 5) and applying a double-validation test to any fitted pre-flare dip candidate (Sect. 4.2).

Figure 3 shows the data median absolute deviation (MAD) on the CHEOPS subarray and detrended CHEOPS imagette light curves, as well as on the flattened 20 s TESS light curves (the CHEOPS subarray data were reduced by its dedicated DRP). The CHEOPS imagette light curves were resampled in 20 s intervals for the sake of comparison: the figure shows that the average MAD of the detrended CHEOPS imagette light curves, even if higher than the MAD of the CHEOPS DRP data, is still lower than the one achieved by TESS on our targets.

4. Flare detection and validation

Flare candidate peaks were searched in the flattened light curves with the PEAKUTILS software (Negri & Vestri 2017), which relies on the inspection of the first derivative of an array, and where we set a minimum peak threshold of four times the noise value. This threshold was chosen following several experiments and visual inspection of the results: it was indeed found that a threshold of 5σ would yield a significant number of missed low-amplitude flares, while a 3σ threshold would produce a large number of false positives, such as clear statistical fluctuations or outliers that were flagged as flares. While most false positives could have been rejected during the flare validation phase, we preferred avoiding the risk of increasing their number. Following visual inspection, we also set a minimum duration of five and three data points for CHEOPS and TESS candidates, respectively, in order to limit the number of false positive detections. If (a) the so-defined flare lasted more than the assigned minimum duration, (b) no part of the flare profile fell in a data gap, and (c) the first two points after the peak candidate were at least 2σ above the noise level, we proceeded with the fit of the flare profile. To do this, we computed a moving median-filtered stellar flux with a window dependent on the data cadence (11 points for TESS data, 29 for CHEOPS); then, we considered the stellar flux from one hour before the time when the median-filtered flux exceeded the noise level, to one hour after the time when the same flux returned to the noise level.

4.1. Flare model and fit

Flares sampled at high cadence often deviate from templates found at lower time resolution, and present multiple peaks, flat-top peaks, Gaussian-shaped bumps and oscillations (e.g. Howard & MacGregor 2022). While analytical models can hardly capture the full complexity of these profiles, they are nonetheless useful to test the accuracy of empirically derived templates, and can be used to validate them and to better identify the quiet stellar flux level.

We modelled each flare profile $f(t)$, t being time, following the double exponential decay that Davenport et al. (2014) found to be representative of most white-light flares observed by Kepler on GJ 1243. This model separates a first, linear decay phase (possibly explained by bremsstrahlung radiation), and a second, gradual phase (probably connected to radiative recombination). Additionally, we smoothed the flare peak following Gryciuk et al. (2017), Jackman (2020), and Mendoza et al. (2022)'s prescription, by taking the convolution of the double exponential with a Gaussian function to represent both an energy release and loss process. This model, which simulates a gradual transition from the rise to the decay phase, proved to be particularly useful to describe the classical exponential shape, the 'bump' and partially the 'flat-top' flare components (e.g. Howard & MacGregor 2022) with the single functional form

$$f(t) = \frac{\sqrt{\pi}AC}{2}\left[(F_1 h(t, B, C, D_1) + F_2 h(t, B, C, D_2)\right], \quad (2)$$

where

$$h(t, B, C, D) = \exp\left[-Dt + \left(\frac{B}{C} + \frac{DC}{2}\right)^2\right] \mathrm{erfc}\left(\frac{B - t}{C} + \frac{DC}{2}\right), \quad (3)$$

and where $\mathrm{erfc}(t) = 1 - \mathrm{erf}(t)$ is the complementary error function. In the above expressions, A is the normalised peak amplitude of the flare, B is the peak time in a rescaled time axis (see below), C its rise timescale, D_1 and D_2 the fast and slow cooling timescales, and F_1 and F_2 (with $F_1 = 1 - F_2$) the relative importance of the two cooling phases.

For each flare candidate, we carried out a least-square fit with the Powell minimisation algorithm implemented in SCIPY (Virtanen et al. 2020), and wrapped in the LMFIT package (Newville et al. 2014). A more thorough exploration of the parameter space was made possible thanks to the basin-hopping optimisation algorithm (Wales & Doye 1997), available in the same library. The method also provides estimates for the parameter uncertainties using the fit covariance matrix.

To improve convergence, we fixed the model parameters to the values Mendoza et al. (2022) fitted on a set of single-peak GJ 1243 flares, and fitted only the flare peak time t_{peak}, its amplitude (which scales the A term) and its full width at half maximum (FWHM). The time axis t in Eq. (2) was then rescaled as $t' = (t - t_{\mathrm{peak}})/\mathrm{FWHM}$[4]. To model any residual trend in the

[4] An example of such an implementation can be found at this GitHub page (https://github.com/lupitatovar/Llamaradas-Estelares)

quiet stellar flux level, a first-order polynomial was included in each flare profile fit.

A final flare validation was performed based on the AIC value of a one-peak flare model against a second order polynomial fit to the flux centred around the candidate peak: if the AIC favoured this flare fit by at least six units, that is, if the model without flares was <5% as probable as the model with one flare to minimise the information loss from the data, the flare was considered validated.

Every validated flare was examined for the presence of additional peaks by adding one flare profile at a time, up to a maximum of five flares. If the AIC of a fit with more flares was preferred to the model with less flares, in a similar way to Davenport et al. (2014), the new fit was retained. Once again, the requirement for the more complex model was set to at least six units a reduction in the AIC value.

To help prevent outliers in complex flare profiles from being misinterpreted as flares, we fixed a lower bound to the fitted flare FWHM for the flare to be validated. We found validating flares with this criterion to be more effective than setting a lower bound to this parameter during the fit, as in the second case the fits of very short-duration flares tended to hit the lower parameter bound. We then required for each candidate flare to have a fitted FWHM larger than the data cadence, that is, 3 and 20 s for the CHEOPS and TESS sample, respectively. Also, to avoid a similar effect in amplitude, we rejected flares with amplitude smaller than 1.1 times the lower bound value.

4.2. Pre-flare flux dips

We used the same method to inspect for the presence of pre-flare flux drops that might be indicative of black-light flares. For each validated outburst, we searched the flux within one hour prior to the rise phase, and fitted a generalised Gaussian model to it:

$$f(t) = \begin{cases} G \exp[-(|t - t_{\rm d}|/w_1)^n] + q & \text{for } t < t_{\rm d} \\ G \exp[-(|t - t_{\rm d}|/w_2)^n] + q & \text{for } t \geq t_{\rm d}, \end{cases} \quad (4)$$

where $G < 0$ is the Gaussian lowest flux value, $t_{\rm d}$ is the Gaussian function central time, w_1 and w_2 represent the Gaussian width on each side, $n > 2$ allows the function to assume less peaked shapes than a standard Gaussian, and q allows for the quiet stellar flux level to be adjusted. We attempted to reduce the dependence on the initial condition for $t_{\rm d}$ by using a basin-hopping optimisation with 10 iterations.

We then compared the AIC of the best fit with the one of a flare without flux drops. As a first validation test against the flare-only model, we retained the dip fit if its AIC was lower than the one without it by at least six units. We also noticed a tendency for the fit to find dip features in data sets with a correlated structure in the quiet stellar flux level, which might be a residual from the detrending or from systematic effects. We exclude stellar granulation as a possible explanation, because its signal is too faint to be noticeable in a single M star TESS (Sulis et al. 2023) or CHEOPS light curve. We then compared the amplitude of the fitted flux drop with the correlated noise level in the light curve after detrending or smoothing. For each CHEOPS target or TESS light curve, this was calculated with Pont et al. (2006)'s method, and the flux drop was compared to the noise level corresponding to the drop's fitted width. We considered a 2.5σ flux drop significance against the red noise level as a threshold for a dip candidate to be visually inspected.

5. Injection tests

To test the accuracy of our detection and fit algorithms, we ran injection tests on the CHEOPS and TESS light curves without detected flares. To create a statistically meaningful sample, we used 100 TESS and 500 CHEOPS light curves, so that some CHEOPS data sets were used for the test more than once. Using the model presented in the previous sections, we added between one and two flares to each CHEOPS visit, and between 10 and 50 flares at random times to each TESS light curve, with log-uniformly distributed amplitudes corresponding to S/N between 1 and 20, and a log-uniformly distributed FWHM between 3 s and 100 min. We also added pre-flare dips before each flare peak, with amplitude between 0.05 and 0.5 times the S/N of the associated flare. The scaling between the S/N of the flare and of the associated dip was adopted to reduce the likelihood of missing large dips following low-amplitude flares. Dips were allowed to have asymmetric widths between 0.1 and 5 min.

We ran our algorithm on the flare-injected light curves after a previous detrending phase, applying the same algorithms described in Sects. 3.1 and 3.2. A simulated flare or dip was considered recovered if the algorithm validated a flare within 3 min before and after its true peak time. Figure 4 presents the recovery rates of flares and dips as a function of their amplitude: the recovery rate was defined as the ratio between the number of validated and injected events in a given amplitude range, and the false positive rate as $n_{\rm false}/(n_{\rm true} + n_{\rm false})$, where $n_{\rm false}$ and $n_{\rm true}$ are the correct and incorrect detection numbers, respectively. In terms of dips, we only show those for which the AIC prefers the model including the flux drop by at least 6 units, and with a significance of at least 2.5 when compared to the correlated noise level of the light curve according to Pont et al. (2006)'s method.

Figure 4 shows that the completeness of our detection algorithm approaches unity for TESS flares with amplitude close to ∼10%, while it is lower for CHEOPS. This can be explained by a combination of poorly corrected systematics and the fact that the largest flares are partly removed by the CHEOPS detrending algorithm, once their duration becomes comparable to the duration of the CHEOPS visits (≃3 h).

Similarly, the completeness of dip detections is higher for TESS than for CHEOPS data. This can once again be explained by both the detrending procedure for the CHEOPS data, which often removes the largest dips, as well as the fact that dips might fall into data gaps.

In Fig. 5, we present the retrieved flare amplitude, FWHM, and dip amplitude against their injected values, for the correct detections. For all pair of retrieved $p_{\rm ret}$ and injected $p_{\rm inj}$ parameters, we fitted quadratic polynomials with the form

$$\log p_{\rm inj} = c_1 \log^2 p_{\rm ret} + c_2 \log p_{\rm ret} + c_3 \quad (5)$$

to the simulated CHEOPS and TESS data separately, and obtained correction factors that we later applied to the results on the observed data. Flare durations, defined as the time during which the fitted flare model is above the photometric noise level, were re-estimated by recomputing each flare profile with the corrected amplitude and FWHM. The c_1, c_2 and c_3 coefficients are presented in Table 1.

6. Results

6.1. Flare rate

We validated 100 and 1364 flares in the CHEOPS and TESS light curves, respectively, considering the components of multi-peak

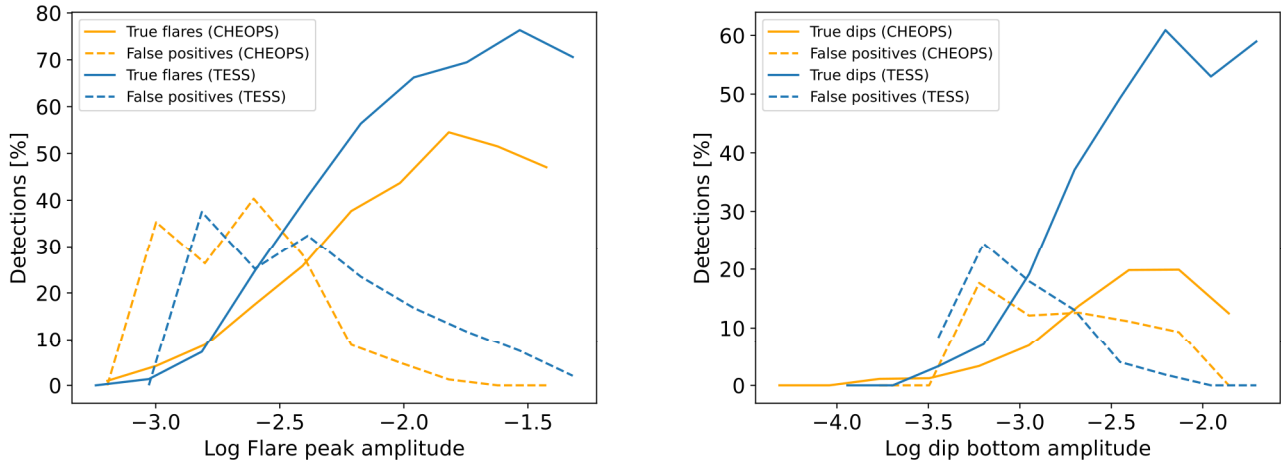

Fig. 4. Correct detections and false positives for flares (left) and pre-flare dips (right). The results for CHEOPS and TESS simulations are shown with different colours.

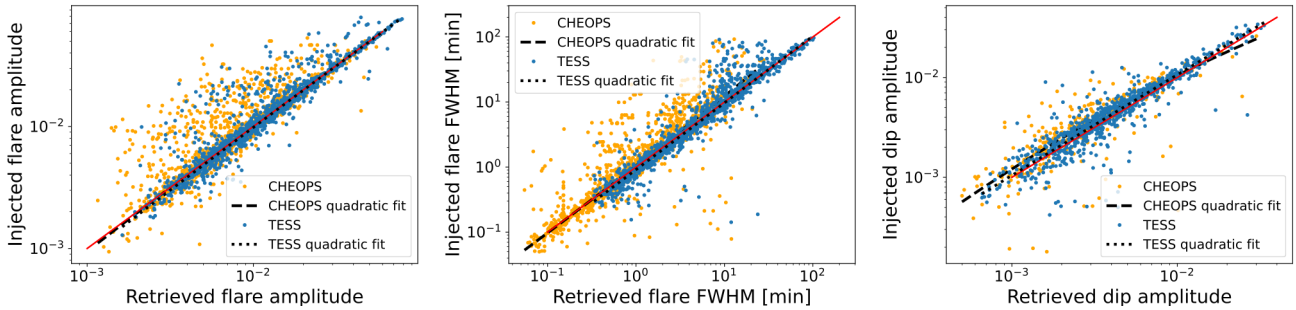

Fig. 5. Retrieved versus input parameters used in the injection tests. The one-to-one relationship is highlighted with red lines, and quadratic polynomial fits for the CHEOPS and TESS data are plotted as dashed and dotted black lines, respectively.

Table 1. Coefficients for the fitted relationship between retrieved and injected flare parameters, for simulated CHEOPS and TESS data.

	Flare amplitude	Flare *FWHM*	Dip amplitude
CHEOPS			
c_1	−0.0302(1)	0.01863(8)	−0.116(2)
c_2	0.877(3)	1.05247(3)	0.36(5)
c_3	−0.120(3)	0.02155(2)	−0.80(7)
TESS			
c_1	−0.01808(7)	−0.00032(3)	−0.0069(4)
c_2	0.947(1)	1.02639(8)	0.973(9)
c_3	−0.041(1)	−0.03846(1)	0.00(1)

Notes. Numbers in parentheses indicate the uncertainty on the last reported digit.

flares as individual events. The left panel of Fig. 6 compares the number of non-flaring and of flaring stars as a function of stellar spectral type. Most detected flares were observed for the latest spectral types, as expected from previous results (e.g. Günther et al. 2020; Jackman et al. 2021).

Comparing targets with both CHEOPS and TESS observations, we found a flare rate in the range 0–15 day^{-1} for objects observed with the former, and between 0 and 4 day^{-1} for those observed with the latter. While these results are broadly compatible with the flare rates found by Günther et al. (2020) on a sample of 24809 M dwarfs observed with 2 min cadence during the first two months of the TESS mission, we refrain from a direct comparison because of the different size of the CHEOPS and TESS validated flare sample.

The validated events here presented are those that passed the fitted duration and amplitude criteria outlined in Sect. 4.1. However, our results might still be affected by false positive detections due to instrumental correlated noise, especially because some of the validated flares have a very low amplitude. Given the results of our injection tests, we inspected the likely contribution of false positives in our results. By reference to the left panel of Fig. 4, we used as thresholds the flare log-amplitude values above which the detection rate becomes higher than the false positive rate: about −2.55 and −2.4 for TESS and CHEOPS, respectively. The flares with fitted amplitudes below these thresholds are 3% for TESS and 35% for CHEOPS. We therefore carried out a visual inspection of the CHEOPS flares validated by our pipeline, to remove obvious false positives: 26 validated flares were rejected. For this reason, the percentage of CHEOPS false positives here presented is much lower than the one obtained from the tests.

6.2. Flare complexity

A combination of Mendoza et al. (2022)'s smoothed, two-phase templates, whose average shape was obtained on a set of single-peaked flares on GJ 1243, was found to successfully reproduce most complex flare components, too. The 3 s cadence data analysis suggests that the higher the cadence, the higher the complexity of flares that might be revealed, with an increase in the fraction of three-peak flares compared to 20 s observations. As shown in the right panel of Fig. 6, about a third of the overall

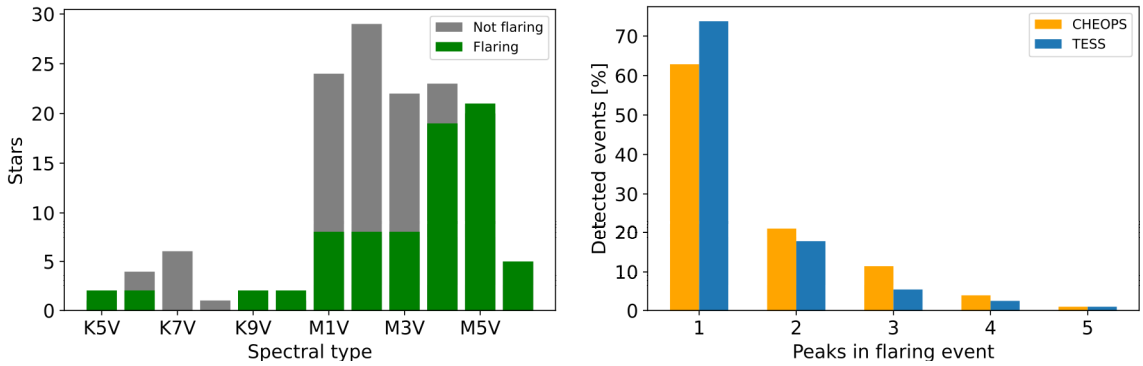

Fig. 6. Flare detection statistics. Left: number of flaring and non-flaring stars as a function of spectral type. Right: percentages of peaks observed per flare event. Flares observed with CHEOPS and TESS are distinguished.

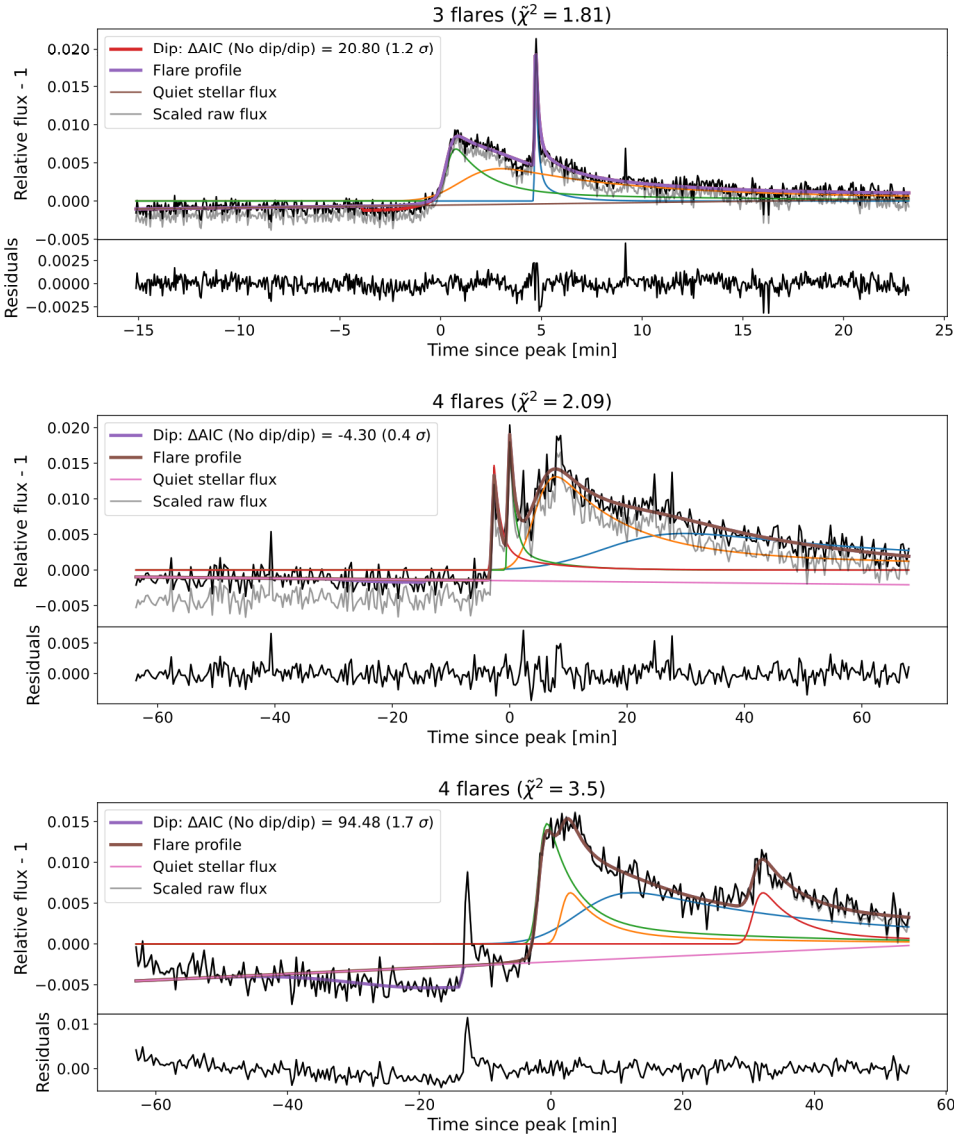

Fig. 7. Examples of complex flare profiles. Top: a bump-peak flare profiles on V1054 Oph. Individual flare components are represented with different colours, and the total model (including the quiet stellar flux) is drawn with a thicker line. The grey, half-transparent line shows the raw flux before light curve detrending. The legend reports the AIC difference between a model without and with dip, and the fitted dip significance with respect to the correlated noise level. The rapid flux drops after min. 15 are likely not to be attributed to the stellar signal. Centre: a flare with structure in the rise phase observed on YZ Cmi. Bottom: a flat-top profile flare on GJ 799B. Our fit strives to model it with a single profile and associates two flares to it. As we found such profiles to be rare, they do not significantly affect our complex flare statistics.

fitted flares presents more than one peak. This confirms the relative importance of complex flares once the observing cadence can resolve their structure (Howard & MacGregor 2022). In particular, many of the two-peak flares are of the peak-bump type discussed by these authors, even if we detected cases where the 'bump' precedes the 'peak'. We also found a few occurrences of structure in the rise phase and a flat-top profiles. Examples for such cases are shown in Fig. 7.

We fitted at most five individual components in multi-peak outbursts, but found by visual inspection flares with an even more complex structure. While fits with a larger number of components could be attempted, care should be exercised against the

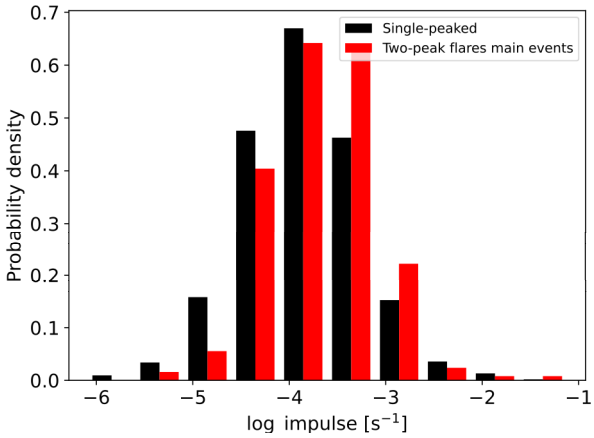

Fig. 8. Impulse for single-peaked flares (black) and the first component of two-peaked flares (red).

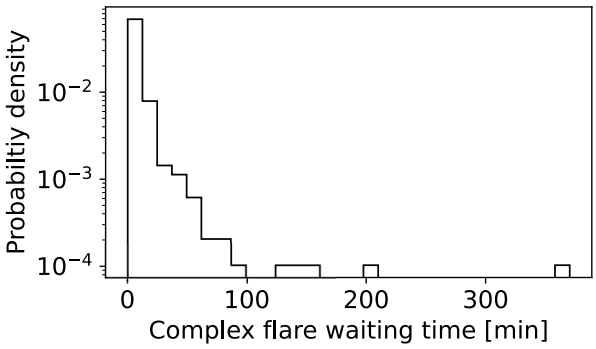

Fig. 9. Distribution of waiting time between consecutive flare components.

increasing number of degeneracies and correlations among the parameters, as well as dependence on the minimisation algorithm. However, our results suggest that particularly complex flares are a minor component in our sample of mostly low-activity stars ($\lesssim$5%). This number accounts also for the QPP candidates that we observed (Sect. 6.7): to be conservative, such candidates were removed when deriving the flare properties discussed in this section. In this regard, we remark that an estimate of a few percent of QPPs in the flare sample is in agreement with the M star literature (Ramsay et al. 2021; Million et al. 2021; Howard & MacGregor 2022).

According to Tovmassian et al. (2003), the peak-bump shape is the fundamental flare shape (even preceded by a starspot-induced pre-flare dip), as energy is re-radiated by the stellar photosphere after the peak phase. However, the complete profile is observable only when the emission site is close to the centre of the visible stellar disc. In this hypothesis, flares that occur near the stellar limb have their 'bump' part less visible, but if they are particularly powerful and their peak occurs in the hidden part of the stellar disc, we can only see a residual 'bump'. As the bump would be visible from the largest fraction of the stellar hemisphere, these authors claim flares with sharp rise and decline but no bump should be about three times less abundant compared to the peak-bump type.

We could not directly test this hypothesis, as our implementation does not clearly separate 'sharp' and 'smooth' single-peak flares. We could, nonetheless, attempt a distinction between these two flare types through the distribution of their impulses, given by the ratio between their peak amplitude and FWHM (e.g. Howard & MacGregor 2022). If Tovmassian et al. (2003)'s model is correct, then the first event in each two-peak flare should have relatively high impulse, as it would correspond to the fast rise-fast decay part of the peak-bump flare prototype. The impulse of single-peaked flares should instead be more evenly distributed between both low and high values, which roughly represent the 'bump-only' and the 'peak-only' type, respectively. Figure 8 shows the distributions of the impulse for these two types of flares: on average, main events in flare pairs have larger impulse values, and the null hypotheses of the two distributions belonging to the same one is rejected by a Kolmogorov-Smirnov (KS) test with p-value of 0.003. However, the distribution of single-peak flares is not evenly distributed among all values: this might be related to the little likelihood of finding a powerful enough flare which leaves its 'bump' imprint right behind the stellar limb.

To conclude, we inspected the waiting time in complex flares, that is, the time between peaks of consecutive individual components (e.g. Hawley et al. 2014). We found most events to happen a few minutes after a preceding outburst, as shown in Fig. 9. This finding recalls solar sympathetic flares, namely pairs of flares which happen in distinct but connected active regions. There, the interval between twin flares can last up to a few hours, depending on the solar cycle phase (Mawad & Moussas 2022).

6.3. Simple and complex flare morphology

Hawley et al. (2014) divided single- and multi-peak flares in the study of their parameters. They found that simple and complex flares have a similar amplitude distribution, but that complex flares tend to be longer and more energetic. However, their result might be affected by the fact that they did not separate the individual components of complex flares. We instead distinguished actual single-peak flares from the components of complex outbursts, while taking care of the fact that the comparison of flare amplitudes obtained for different stars is biased by different stellar magnitudes. Following Raetz et al. (2020), we converted flare amplitudes to flux by adopting the CHEOPS and TESS zero-points and effective wavelengths provided by the filter profile service of the Spanish Virtual Observatory (Rodrigo et al. 2012). By using the stellar distances and magnitudes (Sect. 2), we obtained the luminosities. For CHEOPS data, we used the targets' *Gaia G* magnitudes, and for TESS data the TESS magnitudes available in each data set fits file.

Once flare luminosity is considered and complex outbursts are separated in their individual components, not all parameters of single and multi-peaked events are indeed distributed in the same way, as illustrated in Fig. 10: a KS test on the two luminosity distributions cannot distinguish them with a p-value of 0.86, while the duration distributions are distinguished with $p \sim 10^{-49}$.

It was also previously reported, using *Kepler/K2* data (e.g. Hawley et al. 2014; Raetz et al. 2020), that white-light flare luminosity L and duration d are linearly correlated for flares with duration between ~5 and ~100 min. The coefficients of this linear relationship depend on the time cadence and the noise level of the data, as found on Kepler flares observed simultaneously in long and short cadence (Yang et al. 2018). In particular, a higher cadence allows the resolution of higher peak values, and a lower noise level enables the detection of longer flare durations. The