

CHEOPS Observers Manual



Reference: CHEOPS-UGE-PSO-MAN-001

Issue: 5.0

Status: Released for the 5th Announcement of Opportunity (AO-5)

Authors: David Ehrenreich¹, Mathias Beck¹, Nicolas Billot¹, Christopher Broeg², Magali Deleuil³, Andrea Fortier², Thibault Kuntzer^{1,2}, Kate Isaak⁴, Reiner Rohlfs¹, Maximilian N. Günther⁴, Bruno Merín⁴

Document created on 21 April 2017

Released by the CHEOPS Project Science Office on 12 March 2024

Author affiliations

¹Université de Genève, Observatoire astronomique, Switzerland

²Universität Bern, Physikalisches Institut, Switzerland

³Laboratoire d'Astrophysique de Marseille, CNRS, Université d'Aix-Marseille, France

⁴European Space Agency, European Space Research and Technology Centre, Noordwijk, Netherlands

⁴European Space Agency, European Space Astronomy Centre, Villanueva de la Cañada, Spain

Acknowledgements

The authors would like to thank Bruno Chazelas, Adrien Deline, Luca Fossati, Malcolm Fridlund, Davide Gandolfi, Monika Lendl, Celia Sanchez, Pascal Guterman, Sergio Hoyer, and Sérgio Sousa for careful proofreading, suggestions and comments that have improved this document. They are grateful to the CHEOPS Core Science Team, Board and ESA Project Office for their support.

Cover image credits

Front page: ESA / C. Carreau and *SDO/HMI* image of the transit of Venus.

Last page: ESA / C. Carreau.

Key Documentation and weblinks

AD1	Policies and Procedures document – serving as the reference for Annual Announcements of Opportunity (AOs) and the Discretionary Programme (DP), detailing all relevant information, timelines and requirements.
AD2	Link to the CHEOPS Guest Observers (GO) Programme website: https://cosmos.esa.int/web/cheops-guest-observers-programme
AD3	Link to the registration page for the GO Programme: https://www.cosmos.esa.int/web/cheops-guest-observers-programme/registration Registration is required to access the Phase I Proposal Handling Tool, to download the scheduling feasibility checker tool and to download simulated CHEOPS data.
AD4	Link to the CHEOPS Mission website: https://www.cosmos.esa.int/web/cheops

Acronyms

ADU	Analog-to-Digital Unit
AO	Announcement of Opportunity
AOCS	Attitude and Orbital Control Systems
BEE	Back-End Electronics
BJD _{TT}	Barycentric Julian Date (Terrestrial Time)
CHEOPS	CHaracterising ExOPlanet Satellite
CCD	Charge Coupled Device
DPDD	Data Product Definition Document
DPU	Data Processing Unit
DR2	Data Release 2 (referring to GAIA)
DRP	Data Reduction Pipeline
ESA	European Space Agency
ETC	Exposure Time Calculator
FPM	Focal Plane Module
GO	Guest Observers
GTO	Guaranteed Time Observing/Observers
LEO	Low Earth Orbit
LOS	Line Of Sight
M&C	Monitoring and Characterisation
MJD	Modified Julian Date
MOC	Mission Operations Centre
MTL	Mission Timeline
OBT	On-Board Time
OR	Observation Request
OTA	Optical Telescope Assembly
PI	Principal Investigator
PHT1	Proposal Handling Tool for Phase 1
PHT2	Proposal Handling Tool for Phase 2
ppm	Parts per million
PSF	Point Spread Function
PSO	Project Science Office
px	Pixel(s)
QE	Quantum Efficiency
RMS	Root Mean Square
RON	Read-Out Noise
RTS	Random Telegraphic Signal
SAA	South-Atlantic Anomaly
SED	Spectral Energy Distribution
SEM	Sensor Electronics Module
SFC	Scheduling Feasibility Checker
SNR	Signal-to-Noise Ratio
SOC	Science Operations Centre
SW	Software
TAC	Time Allocation Committee
TM/TC	Telemetry/Telecommands
UTC	Universel Temps Coordonné (coordinated universal time)

KEY DOCUMENTATION AND WEBLINKS	3
ACRONYMS	4
FOREWORD	7
OBSERVING WITH CHEOPS	8
WHAT IS CHEOPS AND HOW DOES IT DIFFER FROM PREVIOUS ESA MISSIONS?	8
STRUCTURE OF CHEOPS OBSERVATIONS	9
<i>Some definitions</i>	<i>9</i>
<i>What are the CHEOPS Science Programmes and how is time shared with the community?</i>	<i>11</i>
PREPARING AND SUBMITTING A PHASE I OBSERVATION PROPOSAL	13
<i>Reserved targets (i.e., lines of sight), unreserved targets, and more</i>	<i>13</i>
<i>Observation feasibility</i>	<i>14</i>
<i>How much CHEOPS time do I need?</i>	<i>17</i>
<i>Contents of a Phase I observing proposal</i>	<i>19</i>
PREPARING A PHASE II OBSERVATION PROPOSAL	19
<i>Contents of a Phase II observation proposal</i>	<i>20</i>
<i>Scheduling Observation Requests (ORs)</i>	<i>20</i>
<i>Detector modes and settings</i>	<i>22</i>
STATUS OF AN OBSERVATION PROGRAMME	23
HOW TO RETRIEVE CHEOPS DATA	24
CHEOPS INSTRUMENT & OPERATIONS	24
THE SPACECRAFT AND ITS INSTRUMENT	24
<i>Overview of the platform and payload</i>	<i>24</i>
<i>CHEOPS in orbit and the observable sky</i>	<i>27</i>
<i>Optical telescope and baffle</i>	<i>30</i>
<i>Focal Plane Module</i>	<i>33</i>
<i>Baseline for the on-board data treatment and data flow</i>	<i>39</i>
<i>Things to keep in mind when preparing your OR</i>	<i>41</i>
PHOTOMETRIC PERFORMANCE	44
<i>Noise budget</i>	<i>44</i>
<i>Example data sets</i>	<i>48</i>
<i>Calibration, monitoring & characterisation of the instrument</i>	<i>52</i>

OVERVIEW OF THE SCIENCE OPERATIONS	58
<i>Mission planning</i>	58
<i>On-ground data processing</i>	59
<i>Data archive</i>	59
CHEOPS DATA.....	59
CHEOPS DATA STRUCTURE.....	59
<i>Data levels</i>	59
<i>Types of CHEOPS files and file nomenclature</i>	60
<i>CHEOPS file structures</i>	62
<i>Headers, keywords and relationship to Phase 2 parameters</i>	63
<i>Errors and data quality flags</i>	63
<i>Coordinate systems</i>	63
<i>Time references</i>	64
DATA REDUCTION PIPELINE (DRP): OVERVIEW AND ALGORITHMS.....	64
<i>Pipeline processing overview</i>	64
<i>Structure of the pipeline (module description & flowcharts)</i>	65
<i>Data flow through the pipeline</i>	66
<i>Data reduction products, report and reference files</i>	75
<i>What the DRP does not correct for (aka, systematic errors)</i>	76
<i>Evolution of the DRP</i>	76

Foreword

The CHEOPS Observers Manual aims to provide members of the community with an easy-to-grasp documentation about CHEOPS observations. Section 0 briefly describes the CHEOPS mission and walks potential users through the different steps of preparing an observation proposal. Section 0 provides more information about the CHEOPS instrument itself and its estimated performances. Section 0 describes CHEOPS data and how they are processed.

This issue of the document is applicable to the 5th Announcement of Opportunity (AO-5) for Observing Time in the ESA CHEOPS Guest Observers (GO) Programme.

Observing with CHEOPS

What is CHEOPS and how does it differ from previous ESA missions?

The European Space Agency's (ESA's) CHaracterising ExOPlanet Satellite (CHEOPS) is the first space mission dedicated to the search for exoplanetary transits through high precision photometry of bright stars already known to host exoplanets. It provides the unique capability of determining accurate and precise radii for a subset of small exoplanets (Neptune-size and smaller) discovered by the new generation of space- and ground-based photometric or spectroscopic surveys. Additionally, it allows mass determinations of exoplanets in resonant orbits via transit timing measurements, and the atmospheric characterisation of highly irradiated exoplanets via occultation and phase curve observations. By focusing on transiting exoplanets with high potential for in-depth characterisation, CHEOPS provides prime targets for future instruments suited to the spectroscopic characterisation of exoplanetary atmospheres, such as NASA/ESA HST, NASA/ESA/CSA JWST, and ESA Ariel.

CHEOPS is an S- or small-class mission, and a partnership between the European Space Agency's (ESA's) Science Programme and the Swiss Space Office. The CHEOPS mission consortium is led by the University of Bern, with important contributions from Industry and Institutes from Austria, Belgium, France, Germany, Hungary, Italy, Portugal, Spain, Sweden and the United Kingdom. The responsibilities within the mission are quite different from traditional medium- or large-class ESA missions. Specifically, ESA is the CHEOPS mission architect and is also responsible for procurement of the spacecraft and for the Guest Observers (GO) Programme. The Consortium is responsible for providing the science payload, mission and science operations, the science team and performance monitoring and evaluation.

The CHEOPS payload consists of a single instrument – a space telescope with an effective diameter of 30 cm, which feeds a single CCD focal plane detector covering 330–1100 nm with a field of view of 0.32 deg². The payload design and operation have been optimised to achieve ultra-high photometric stability, achieving a photometric precision of 20 parts per million (ppm) on observations of a G5 dwarf star in 6 hours, and 85 ppm observations of a K5 type dwarf star in 3 hours. In both cases the ultra-high stability shall be maintained over a minimum period of 48 hours. The point spread function (PSF) of the payload is approximately 16.5" (90% encircled energy radius). These characteristics, together with the high-cadence (better than one-minute sampling), offer the capability to measure high-precision, high-cadence, broad-band light curves covering the visible-near infrared waveband.

CHEOPS was launched on 18 Dec 2019 and, after successful in-orbit commissioning, the science operation phase formally started after successful In-orbit commissioning review on 25 Mar 2020. The nominal mission (NM) duration of the science operations phase is 3.5 years, i.e., running from 25 Mar 2020 until 24 Sep 2023. ESA's Science Programme Committee meeting on 7 Mar

2023 approved a first mission extension (EM1) from 25 Sep 2023 until 31 Dec 2026, and an indicative second mission extension (EM2) from 1 Jan 2027 until 31 Dec 2029. During the NM, twenty percent (20%) and during EM1 thirty percent (30%) of the science time on CHEOPS are made available to the astronomical community through the GO Programme that is open to scientists world-wide. Details of the programme can be found at the link below:

→ <https://www.cosmos.esa.int/web/cheops-guest-observers-programme>

In contrast to Medium- and Large-class missions, which are operated by ESA from the European Space Operations Centre (ESOC, Germany), CHEOPS is operated by its mission consortium: communications with the satellite are sent and received at the Mission Operations Centre (MOC) located in Torrejón, near Madrid (Spain). The telemetry packets (Level-0 data) are transmitted to the Science Operations Centre (SOC) at the astronomical observatory of Université de Genève, Switzerland. There, the data, received as telemetry packets, are pre-processed (Level-0.5 data), reduced into images (Level-1 data) and light curves (Level-2 data) and archived¹. The SOC also performs the mission planning and scheduling of all CHEOPS observations. The scheduling instructions are transmitted as telecommands to the MOC, from where they are uplinked to the satellite. The scientific exploitation of the satellite is shared between the mission consortium and ESA for the Guaranteed Time Observation (GTO) and GO Programmes, respectively (see §0). The Project Science Office (PSO) coordinates these aspects by following-up CHEOPS observations, performing instrument performance monitoring, issuing and maintaining the relevant tools and documentations to observe with CHEOPS.

Structure of CHEOPS observations

Some definitions

CHEOPS conducts scientific observations from its **orbit** around the Earth. These observations are structured with **programmes**, **observation requests**, and **visits** (see Figure 2). Due to orbital constraints, these observations can often be **interrupted**, leading to the notion of **observing efficiency**. One of the big challenges of the mission is to perform many **time-critical observations** in these conditions. The bold-face terms above will be used throughout this manual and should be understood to facilitate further reading. They are defined below.

Orbit — CHEOPS revolves around the Earth in Sun-synchronous, low-Earth orbit (700 km altitude). The spacecraft completes one orbit around the Earth in 99 min. This duration is used as the time unit for visits (see below). The science time of an orbit is usually shorter than the orbit duration due to interruptions by the South Atlantic Anomaly (SAA; 0% to 20% of an orbit) or Earth occultations.

¹ See section 3 for further details on the CHEOPS data structures and level data products.

Interruptions — It can be the case that fractions of a CHEOPS orbit cannot be used for science observation. This can happen when the target is hidden by the Earth during an Earth occultation (which happens for certain pointing directions at certain times of year), or barely visible (with substantially degraded photometric precision) due to stray light from the illuminated Earth limb or particle hits during passages through the SAA. Data are normally not downloaded during interruptions. More details on interruptions are provided in §0.

Visit — A visit is a sequence of successive CHEOPS orbits devoted to observing a given target. Visits are the building blocks of CHEOPS observations. See an example in Figure 1.

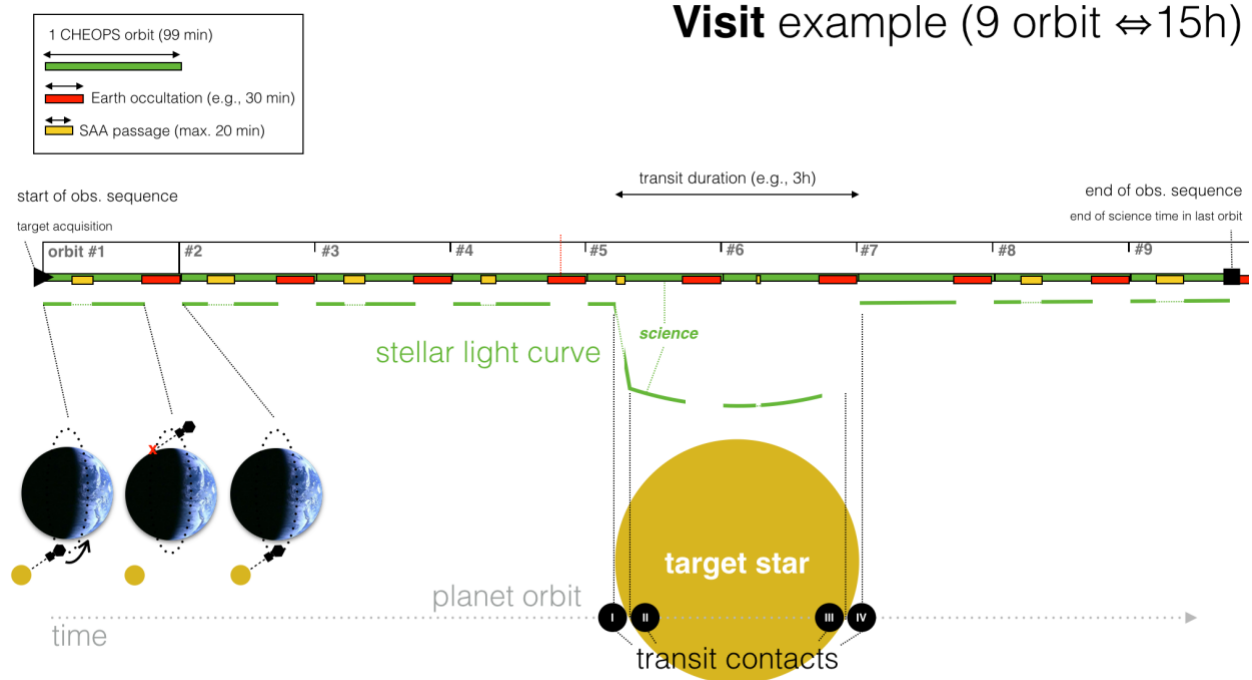


Figure 1: Example of a CHEOPS visit.

Observing efficiency — The observing efficiency is the ratio between the amount of science observing time available during a visit (excluding the interruptions) and the total amount of time in a visit (including interruptions).

Observation Request (OR) — A set of one or more identical visits aiming at a given target. As an example, achieving a given signal-to-noise ratio on an exoplanet transit depth by co-adding several transits can be typically dealt with within a single OR.

Time-critical Observation Request — An OR that must be timed to observe a transient or periodic astronomical event. CHEOPS has been designed to observe predominantly planetary transits, which are periodic. Transit ephemerides must be provided within the OR for such observations.

Programme — A (science) programme is an ensemble of ORs. It has a title, a Principal Investigator (PI) and an identifying number. Note that the GTO Programme and the GO Programme each consist of a set of science programmes.

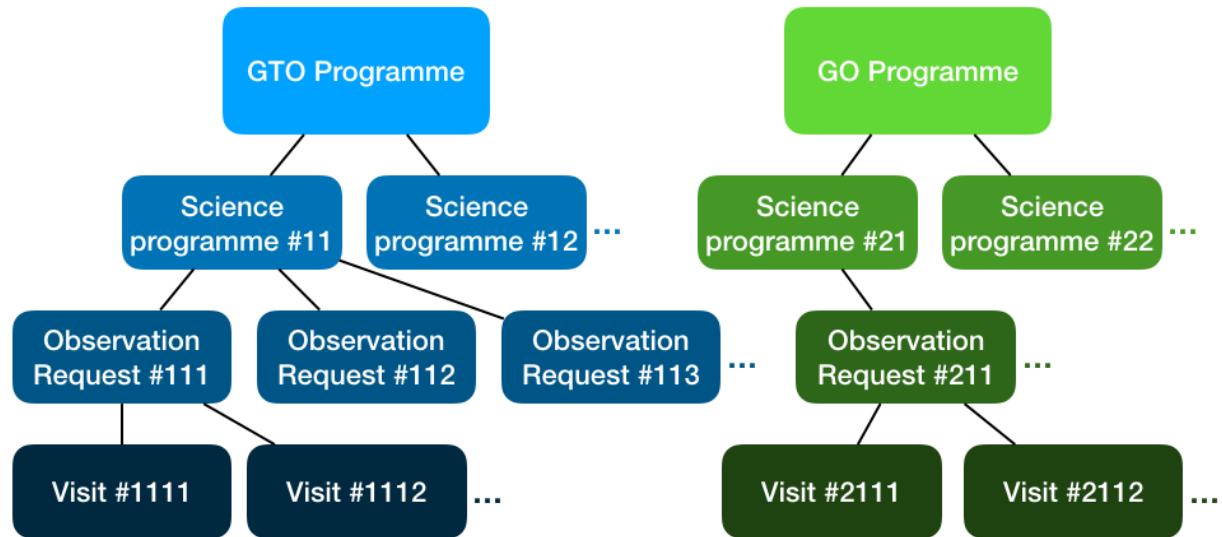


Figure 2: Sketch of the organisational structure of observing programmes on CHEOPS.

What are the CHEOPS Science Programmes and how is time shared with the community?

Science observing time on CHEOPS is split between a GTO Programme, run by the CHEOPS Mission Consortium, and a GO Programme, run by ESA.

Up to ten percent (10%) of the total time on the mission is top-sliced for engineering (5% covering e.g. instrument and platform maintenance, safe mode, etc.) and for characterisation and monitoring activities (see §1.1.1.20).

The CHEOPS Mission Consortium (representing ~40 institutions spread over 11 ESA member countries) has built the CHEOPS instrument and operates the satellite. It has set up a Science Team to establish the GTO Programme. Of the nominal science operations time, at least eighty percent (80%) in the NM and seventy percent (70%) in the EM1 are available to the GTO Programme. The Science Team includes five ESA-appointed members in the NM. It includes the ESA PS as an ex-officio member, five ESA-appointed members, and six respective Collaborators (one per member) in the EM1.

In the NM, the GTO programme covered six broad scientific themes in Exoplanet science:

- Finding transits of known exoplanets;
- Improving radii of known transiting exoplanets;

- Exploring systems in search for new exoplanets;
- Characterising exoplanet atmospheres;
- Uncovering new exoplanetary features;
- Stellar and planetary science relevant to the analysis/interpretation of exoplanet data.

In the EM1, the GTO will focus on three Axes of maximum scientific return:

- Architecture of systems and structure of planets (Axis 1)
- Atmospheres and climates of irradiated planets (Axis 2)
- Deep characterisation of planets and their environments at the frontier of exoplanet science (Axis 3)

A share of the CHEOPS GTO ($\leq 15\%$) will be dedicated to building community synergies through the Synergy with other Missions (SoM) programme run by the Consortium.

A description of the GTO programme at the time of the opening of the various AOs can be found at this link:

→ <https://www.cosmos.esa.int/web/cheops/the-cheops-guaranteed-time-observing-programme>

This set of programs and their associated targets are expected to evolve during the mission. ESA runs the GO Programme for the community. Of the nominal science observing time, up to twenty percent (20%) in the NM and up to thirty percent (30%) in the EM1 is dedicated to the GO Programme. In each case, this is split between annual Announcements of Opportunity (AOs) and the Discretionary Programme (DP) which runs throughout the year.

As of the EM1, the GTO Programme maintains a smaller Reserved Target List (RTL) of only up to 50 targets. See §**Error! Reference source not found.** for details.

Scientists can submit proposals to observe with CHEOPS regardless of affiliation. All proposals must be submitted to ESA via a two-step process, and according to the guidelines and policies relevant to the Call (e.g., issues of AD1 for AOs and DP).

In Phase I, observers are requested to provide inputs that enable both a scientific and technical evaluation of the proposed observations to be made.

As of EM1, this first step is based on a double-anonymous peer-review by an independent Time Allocation Committee (TAC). Please carefully consult the Policies & Procedures document and consult with the ESA PS (cheops-support@cosmos.esa.int) if you have any questions.

Following the award of telescope time, PIs will be requested to provide inputs to Phase II, which include detailed information relevant to scheduling, planning and execution of the observations.

Users should note that a valid ESA Cosmos ID is needed to submit a proposal and access at least one of the tools needed to prepare proposals – no ID/permission, no submission! ESA Cosmos is the portal for users of ESA's science missions and many users of ESA missions will already have such credentials: please request an ID and register to use the CHEOPS GO website as soon as you can, not just before the closing date of the Call. It is important to note that the

same user credentials (user ID) will be used to generate the account details needed to complete Phase II inputs and to access proprietary data in the CHEOPS Mission Archive.

Users need a valid Cosmos ID to submit a proposal. A Cosmos account and the corresponding ID can be queried at the link below:

→ <https://www.cosmos.esa.int/web/cheops-guest-observers-programme/register>

Preparing and submitting a Phase I observation proposal

After having decided that the top-level capabilities of CHEOPS match your needs, it is time to determine whether CHEOPS really is useful to you – specifically to confirm that your targets are not already reserved by other programmes, that they are visible to CHEOPS, and to determine how long it will take to achieve the necessary sensitivity.

Reserved targets (i.e., lines of sight), unreserved targets, and more

The CHEOPS Mission Consortium provides a list of targets that are reserved for the GTO programme, called the Reserved Target List (RTL). Targets on the RTL are not open for new GO proposals. It is therefore important that users check in advance that their targets are not already on this list.

As of the EM1, the GTO Programme's RTL can contain up to 50 targets. All other targets are unreserved and open for GO proposals. However, some special cases may apply (e.g., if the target is already part of an active GTO programme but unreserved; or if a target is part of an active GO programme). Please carefully consult the Policies & Procedures document and consult with the ESA PS (cheops-support@cosmos.esa.int) if you have any questions.

At any time, the status of a given target (e.g., GTO reserved, or part of active GTO or GO programmes) can be checked using the Target and Programme Information and its Target Checker available via:

→ <https://cheops.unige.ch/pht2/search-reserved-targets/>

The target name can be resolved with SIMBAD, or the target coordinates can be provided directly to query the list. **Some target denominations are not referenced in SIMBAD** (e.g., TICs or TOIs); in case of doubt, it is always safer to use target coordinates. Evidence that the RTL checker has been used to verify that a target is not reserved will need to be submitted as part of the Phase I proposal submission process (see AD1 for further details).

For convenience, a csv file containing the RTL at the time of the opening of an annual Announcement of Opportunity for the GO Programme can be consulted on the ESA GO website (see AD2).

Both the GTO and GO components of the target list will evolve after annual calls, and proposers should always use the Search Reserved Targets tool to check whether a target is reserved (<https://cheops.unige.ch/pht2/search-reserved-targets/>).²

Observation feasibility

The first thing to check is whether your target is visible for CHEOPS. From its near-polar, Sun-synchronous orbit, CHEOPS can access a large fraction of the sky and has the best visibility in the ecliptic plane. It has, however, a few limiting pointing constraints, which are detailed and illustrated in §0. The visibility decreases towards the ecliptic poles. The spacecraft has a stringent 120-degree Sun-exclusion angle. This implies that objects having ecliptic latitudes at the epoch of observation above 60 and below -60 degrees cannot be observed (Note that the latitude and longitude of celestial objects can oscillate by up to 20" with 18.6-year period due to the nutation of the Earth axis). It is therefore feasible to easily rule out the possibility of observations of targets that are close to the ecliptic poles.

1.1.1.1 Visibility maps cumulated over a year

The map of the sky visibility, cumulated over a year, is shown in Figure 3. It can be used to make a first assessment of whether your target is visible with CHEOPS.

Note about the Kepler field — The Kepler field lies at the limit of the CHEOPS visibility zone. It should generally be understood that Kepler targets are mostly out of reach of CHEOPS. However, there are exceptions towards the extreme south of the field (see Figure 4 for a close-up).

A first-glance assessment of a target's visibility can be achieved using the detailed sky visibility maps available on the ESA mission website (see AD4).

Stray light (from Earth) dominates the photometric precision noise budget for magnitudes $V \approx 11$ and fainter (see §0 and Figure 19). Because the stray light contamination depends on where the telescope is pointing and where the telescope is in its orbit, the amount of time during which a faint target can be observed with a given photometric precision can be reduced. In other words, there are interruptions due to stray light, during which the target remains visible, but the photometric precision can be expected to decrease. For this reason, the detailed visibility maps cumulated over a year are provided (i) for stars brighter than $V=9$ and (ii) for stars with $9 < V \leq 12$. Interruptions can also be caused by Earth occultations and SAA crossing, so that only a part of the orbit duration can be used for observations. The maps are thus also calculated assuming different fractions of time available during an orbit.

² The target list could also evolve during the time that the Call is open in case of the acceptance of a DP proposal. It is thus wise to check the Search Reserved Targets tool in all cases.




Figure 3: Annual sky visibility map. Colour code shows the total number of days a given position in the sky can be observed during a year, allowing for 50% orbital interruptions. White regions are excluded. The black line represents the ecliptic plane. Because of its low-Earth orbit, most CHEOPS observations will be interrupted. This map allows for interruptions of up to 50%. The fraction of the visible sky available with low interruptions ($\leq 20\%$) is smaller, see Figure 11.




Figure 4: Detailed visibility of the Kepler field. The Kepler field limits are given by the red line. Most of the Kepler field cannot be observed; however, a small portion in the extreme south is visible with CHEOPS for a limited time each year.

These maps can only provide best-case scenarios for targets fainter than $V=12$, which is the nominal faint magnitude limit of CHEOPS. In-orbit commissioning has shown that CHEOPS works best for bright stars. For stars fainter than $V \approx 11$, the photometric performance varies a lot depending on stray light and background contamination. Users proposing targets at the faint limit or fainter may expect the photometric precision of the light curve products to vary, especially depending on field crowding (low field crowding is an advantage). For such cases, a higher level of photometric detrending (not included in the Data Reduction Pipeline; DRP) should become necessary to exploit the data. As faint-target observations may be carried out under conditions that could be detrimental to the scientific goal of the observations, proposals with such faint objects will be dealt with at the discretion of the TAC.

1.1.1.2 Visibility maps cumulated over a month

Because the orbit of CHEOPS precesses in 1 year (remaining perpendicular to the direction vector of the Sun), the telescope sees different patches of the sky along the year. The instantaneous visibility of CHEOPS is approximately similar to that of a ground-based telescope located close to the equator at the same date. Monthly visibility maps (together with the data used to generate them) can be found on the ESA GO website (AD2) and are shown in Figure 5.




Figure 5: Instantaneous (monthly) visibility maps for CHEOPS.

1.1.1.3 Visibility Tool (recommended for Phase 1)

The new Visibility Tool, available as a Python script and Jupyter notebook, allows you to swiftly and easily check whether and when a target is visible to the instrument. You can also estimate the average observing efficiency for your target over the course of the year. The tool is available at:

→ https://gitlab.unige.ch/cheops/CHEOPS_visibility_tool

1.1.1.4 Scheduling Feasibility Checker (recommended for Phase 2)

Photometric precision or timing precision, or a combination of both, could be critical to some science programmes. For these programmes, it is important to know in advance if, given the visibility constraints and orbit interruptions, the required precision(s) can be achieved. The best assessment of the visibility can be obtained with the Scheduling Feasibility Checker tool, available with its manual at:

→ <https://gitlab.unige.ch/cheops/containers/cheops-sfc-webtop>

How much CHEOPS time do I need?

After checking the target visibility, the next step is to determine how much time shall be requested to achieve the required photometric precision. The photometric precision achievable in a given time interval can be calculated using the **Exposure Time Calculator (ETC)**. The ETC together with its manual can be found at:

→ <https://cheops.unige.ch/pht2/exposure-time-calculator>

Use of the Exposure Time Calculator is mandatory for all Phase 1 proposals. A single pdf file generated with the ETC “export results” functionality for each target in the proposal will need to be provided at Phase I submission.

Please keep in mind that only a fraction of this time interval will be spent integrating the target photons due to interruptions. The ETC allows the users to specify a given observing efficiency – or estimate it depending on the target coordinates. Furthermore, the user can input a flux value in electrons per second, or let it be estimated through the *Gaia* (*G*) magnitude. The ETC takes as input the following parameters,

- star spectral type;
- star *G* magnitude from *Gaia*;
- star effective temperature;
- duration of the observation;
- target coordinates (right ascension, declination);
- exposure time.

The main output is the expected level of noise on the photometric measurement for the duration of the observations, and for several pre-specified time bins.

Note that the user must specify the exposure time; that is, the time during which photons are collected to record an image (minimum allowed value is 0.001 s, maximum allowed value is 60 s). The first rule to select the exposure time is to make sure that no pixel inside the PSF saturates. This can be checked with one of the outputs of the ETC, the "Saturation level at PSF peak [% of FWC]". For the selected exposure time, the ETC calculates the percentage of the full well capacity of a pixel filled by the highest peak of the PSF. The user should always verify that it is below 100%. Actually, the *maximum suggested exposure time* is the one where the PSF peak fills 85% of the full well capacity (to provide enough margin to clearly avoid saturation). It is also recommended to select an exposure time close to this maximum suggested exposure time (or 60s if the 85% cannot be achieved) to keep the instrumental noise to a minimum.

On the other end, the *minimum suggested exposure time* is the one corresponding to the PSF peak filling 10% of the full well capacity.

The user is responsible for selecting the correct exposure time and may have reasons not to follow the guidelines given above. For example, the user might be interested in shorter exposure times than recommended to have a faster cadence of images or imagerettes (see Table 2).

Please note that each OR accepts only one exposure time setting, meaning that all exposures of all visits in an OR will be identical.

For each target, the total time to request is the duration of the visit, possibly multiplied by the number of visits (with identical settings) in case several observation epochs are required. If visits with different settings (e.g., a different visit duration) are required in the proposal, then the total time requested will be the addition of all visits. The total time request for the proposal will be the sum of the total time requested for each target.

CHEOPS is designed to achieve excellent *relative* photometric precision (especially important for exoplanet transits). Therefore, observers are generally advised to reserve enough time to acquire a proper baseline before or after (or both) the transient event they are interested in.

The duration of a single visit must be comprised between 1 CHEOPS orbit (99 min) and 100 CHEOPS orbits (approximately 7 days).

Note, however, that CHEOPS is designed to maintain its photometric performances for periods of time up to 48 hours (including interruptions). Early operations of the satellite have shown that there is no sudden degradation of the photometric performance after 48 hours; however, users must be aware that the photometric precision of CHEOPS is not guaranteed beyond this time scale. Observations lasting longer than 100 CHEOPS orbits can be arranged for very rare exceptions but must be broken in several adjacent visits.

Contents of a Phase I observing proposal

A Phase I CHEOPS observing proposal is made up of three components:

- Scientific justification
- Technical justification and implementation
- Input to the on-line proposal submission tool

Further details on each of the different components, as well as how and when to submit proposals can be found by going to the link below:

→ <https://www.cosmos.esa.int/web/cheops-guest-observers-programme/>

and looking at the webpage dedicated to the particular AO Call or DP. Note that the only way in which proposals can be submitted is via the online proposal submission tool which can be found by following the links given on the Call webpage.

All Phase I proposals will need to be completed and submitted on-line by the AO deadline specified in the AO document and given on the webpage for the Call.

Preparing a Phase II observation proposal

The PIs of proposals that have been awarded time will be requested to follow the second phase of the proposal preparation. The activity will start shortly after the announcement of successful Phase I observation proposals is made and will be open for the duration noted in the Call timeline. The purpose of Phase II is to provide the detailed and accurate information required to schedule observations, in the form of ORs. Target information (including coordinates) together with the time awarded by the TAC for each target will be automatically carried over to the **Proposal Handling Tool for Phase II (PHT2)** and will not be modifiable by the user. Access to the tool is provided by the Project Science Office at

→ <https://cheops.unige.ch/pht2>

Access will require user credentials, which will be provided to the PIs of successful Phase I proposals shortly after they are informed of their award of time.

ORs submitted for Phase II will be consolidated by the Project Science Office. The ORs from the GO Programme together with those from the GTO programme will be used as input to an automated observation scheduling tool, which will generate a long-term observing schedule for science operations covered by the period of the Call.

The execution of a GO programme accepted in Phase I is pending timely Phase II submission and is on the assumption that there are no issues with the submitted ORs. It is very important to note that even GO programmes that have passed through Phase II validation may end up not being executed if they have too strong timing constraints, if their target locations are challenging, if their duration is especially long, or if they include

any combination of these criteria. Finally, scheduling of a given OR may be made difficult due to competition from ORs in other programmes with similar constraints.

Contents of a Phase II observation proposal

CHEOPS observation programmes are built up from individual ORs. One OR contain one or more identical visits (same target, same duration, etc.). Important information requested at Phase II includes:

- Target coordinates, proper motions, *G* magnitude, effective temperature (for stars), parallax and GAIA DR2 ID.
- Time-critical constraints, i.e. for planetary transits, transit ephemerides (transit mid-time epoch, revolution period and a start phase slack), earliest start and latest end dates.
- Detector settings (such as exposure time).

Users are responsible for the accuracy of the parameters they enter in Phase II.

Scheduling Observation Requests (ORs)

There are two types of nominal science ORs: **Time-critical** ORs are designed for observing transient events at specific epochs or periodic events, typically planetary transits. **Non-time-critical** ORs are appropriate for other types of targets, i.e. those that do not require precise timing to observe an event at a particular time.

1.1.1.5 Observing efficiency

Interruptions of orbital nature create gaps in CHEOPS photometric time series. These interruptions can be predicted (to some extent), which means that in theory, one could minimise the effect of interruptions by tailoring the observation schedule. The automated Mission Planning Scheduler will perform this optimisation, but it should be clear that ***there will be interruptions in most CHEOPS observations***. The fraction of science time in the visits of an OR is called the **observing efficiency**. A requested observing efficiency of 60% means that at least 60% of the orbit (see limitation below) will be collecting photons from the CHEOPS target and that the observations could bear as much as 40% interruptions.

The observing efficiency of a given target is mainly driven by its position on the sky and the spacecraft orbit. For example, winter targets with high declinations have intrinsically lower observing efficiencies than summer targets around the ecliptic. To second order, the observing efficiency is also affected by the date of the observation. For example, observations taken at a time when the target is close to the anti-Sun position will have higher observing efficiencies than if it is observed while it sits at the edge of the seasonal visibility window. This timing aspect is optimised by the schedule solver algorithm to find observation dates that maximise the overall observing efficiency.

The general recommendation is to request an observing efficiency of 50% as a default value (average efficiency for winter targets is only around 55%).

If a science case requires a higher efficiency, users must use the **Scheduling Feasibility Checker** to assess its impact on the schedulability of their observations (see §1.1.1.3). Alternatively, users may also decide to set up their observations with a modest global observing efficiency but with a higher local efficiency at specific phases of a transit observation, typically around transit egresses.

In any case, please keep in mind that requesting high observing efficiencies or local critical phase ranges does reduce, sometimes drastically, the schedulability of an observation.

Observers may also request some level of flexibility on the start time of their observations. Such flexibility allows the schedule solver algorithm to shift visits during the optimisation process to improve the filling factors of the overall schedule and, therefore, to schedule more observations. However, such flexibility comes at a price: for visits requesting scheduling flexibility, particularly those shorter than 3 orbits, the effective observing efficiency may end up being lower than the requested value, sometimes by up to ~15%. This may happen under special circumstances, typically when the schedule solver algorithm adjusts the visit start time to optimise the overall schedule, which may result in a visit being shifted toward the SAA, Earth occultations or straylight regions.

Due to the competition from other time-critical programmes, users should be extremely cautious about requesting very high observing efficiencies, as this could cause a programme to become non-schedulable. Don't be greedy!

1.1.1.6 Setting time constraints

Time constraints such as the earliest start and latest end dates at which an observation should be executed can be provided at the level of ORs, not on visits. If several visits on the same target have to be scheduled at different epochs, then several ORs must be created with their own time boundaries. These constraints must be provided as Barycentric Julian Dates, Terrestrial Time (BJD_{TT}).

1.1.1.7 Setting the ephemeris of a periodic event

Any periodic event can be located in time, provided a reference date/time and the period are known. A convenient parameter used to measure the progress of a periodic event is the phase, which is the time divided by the event period. By definition, the reference time (e.g. the mid-transit time) is at phase 0. In the Proposal Handling Tool Phase II, valid values for the phase are between 0 and 1. If in doubt about setting phase constraints on a periodic event, one should use the Scheduling Feasibility Checker to test and validate observing strategies.

To observe a planetary transit it is, of course, not necessary to observe the whole orbit but just the interesting bit, i.e. the transit itself and an out-of-transit baseline, before and after the transit, for reference (CHEOPS is designed to achieve precise *relative* photometry!). Users can define

when to start (with a margin) using the so-called start phase slack: an interval in orbital phase during which the visit shall start. The longer the slack, the easier the visit will be to schedule. Because of orbital interruptions, gaps in the light curve could mean that CHEOPS will miss a critical portion of the transit, e.g. its ingress. To prevent this, users are allowed to set an additional critical phase range, during which they could define a specific observing efficiency. Be warned, though, that setting a critical phase range places a significant constraint on scheduling and may make the programme un-schedulable. Figure 6 illustrates what the ephemeris parameters are.




Figure 6: Time constraint parameters on top of a sketch transit light curve. The time (T) reference T_c is the mid-transit time, P is the period of the planet and ϕ is the phase: $\phi = (T - T_c)/P - E[(T - T_c)/P]$, where $E[x]$ is the integer part of x . The start phase slack is shown in green and a critical phase range is shown in orange.

Detector modes and settings

Many of the detector settings, such as the read-out mode, will be pre-set as they depend essentially on the target magnitude.

The exposure time is determined by the user, however, following guidelines provided in the Exposure Time Calculator.

The observation cadence is fixed to optimise the use of the (limited) downlink bandwidth. The cadence determines whether it is possible to retrieve individual exposures or only co-added exposures to save bandwidth. More details on the detector modes can be found in §0 and Table 1. The cadence of science images downloaded to the ground can be found as a function of the exposure time in Table 2.

Status of an observation programme

The status of observation programmes can be consulted in the “Programmes” tab of the Proposal Handling Tool for Phase II (PHT2). The different statuses are described below in chronological order.

Draft—This is the initial state of a programme after its creation. At this stage, the programme editor can add new ORs.

Submitted— The programme editor must submit the programme by clicking the 'paper plane' icon. This submission stage is mandatory for SOC to consider the programme complete and ready for the observation planning process. **Programmes in the “submitted” state cannot be edited, nor its associated ORs, until they have been validated by the Project Science Office.** This freezing phase is typically shorter than one week, and it is necessary for a consistent and robust validation process.

Pending Validation—Every Monday morning, the Project Science Office will export all newly submitted programmes, if any, for inspection and technical validation of its ORs. At this stage, the programme status will change from “Submitted” to “Pending validation”. Programmes in the “Pending validation” state cannot be edited, nor its associated ORs, until it has been validated.

Running—All newly submitted and valid ORs are ingested into the Mission Planning System. At this stage, the programme status will change from “Pending validation” to “Running”, meaning that these ORs have become part of the pool of valid candidate requests from which the automatic scheduling algorithm can select.

Disabled—When a "Running" OR is disabled in the Mission Planning System, its status turns to "Disabled" in the PHT2. ORs are disabled in the MPS when they are completed, when their *latest end date* is in the past, or more generally when they are removed from the pool of active ORs to be scheduled. A disabled OR can be cloned for convenience.

Aborted—When an OR is stopped in the PHT2, this information is pushed to the Mission Planning System to disable this specific OR. The OR status then becomes “Aborted”. An aborted OR can be cloned, for convenience.

The PI will receive a notification that a visit has been scheduled typically a couple of days to a week prior to the actual observation. This notification indicates that the telecommands associated with the observation have been uplinked successfully to the spacecraft.

Update of submitted ORs that are part of GO Programmes is not possible by the PIs of the programmes, but requires the intervention of the SOC. This intervention is possible under exceptional circumstances, e.g., when an update to the target ephemeris is required. In such cases the PI should make the request for an update by sending an email to the ESA CHEOPS Project Scientist (cheops-support@cosmos.esa.int), detailing the nature of the update(s) required and the justification for them. All emails to the support desk should include the specific AO Call / DP identifier and programme ID in their subject (e.g. “AO-5: PR250001” or “DP: PR450001”). Details on the steps required to update the OR will then be sent to the PI.

How to retrieve CHEOPS data

Once the observation for a given visit has been made, the data received on ground and processed by the Science Operations Centre in Geneva, the data will be made available via the CHEOPS archive browser. The PIs of the programme from which the visits originate will be notified by e-mail that data can be downloaded. The time from the end of the observation to notification is foreseen to be about three days. **The CHEOPS archive browser is the unique interface to access CHEOPS data** for visits that belong to a GO or DP programme. Science data taken with CHEOPS is subject to a proprietary period (see the Call document for details). Only the PI and the Co-I (if details are provided) listed in the Phase I proposal will be able to access the data – both will need credentials to log in. These credentials will be provided shortly after the announcement of the awards of CHEOPS observing time and are the same that are needed to access PHT2.

CHEOPS instrument & operations

This chapter briefly describes the CHEOPS instrument and platform design. This includes the orbit and visibility constraints and also the camera design. The different instrument modes are briefly outlined and the science data product, which is downlinked to ground, is described. In the second part the expected performances are given in form of the noise budget.

The spacecraft and its instrument

Overview of the platform and payload

CHEOPS has a single scientific payload on a small satellite platform provided by Airbus Defence and Space. The payload is a single instrument: a 30 cm effective aperture reflecting telescope designed to observe individual stars. It is a fully cold redundant system (except for the CCD and CCD clock driver). The total mass of the satellite (payload and platform) is 280 kg, the mass of the payload alone being 60 kg. The spacecraft will provide 60 W continuous power for the instrument operations. Figure 7 shows a view of the instrument integrated on the platform and Figure 8 a front and rear view of the instrument, with a cross-section of the instrument shown in Figure 12.

The main design drivers of the satellite were: to achieve high photometric stability and high thermal stability of the instrument, and to minimize the Earth straylight contamination. At the same time, the aim was also to maximize the observable sky while keeping the radiation dose as low as possible. The major requirement is the photometric stability; therefore, the detector gain has to be extremely stable, and Earth stray light must be suppressed to a very high degree.

In addition to a thermally stable environment, the instrument requires high pointing stability: the telescope line-of-sight must remain stable to 4 arcsec rms (root mean square) within nominal operational conditions³. This precision can be achieved on a small platform by including the instrument data in the attitude control loop.

An S-band system will be used for telemetry and telecommands (TM/TC) downlink and uplink between the satellite and the Earth, allowing data downlink rate of 1.2 Gbit/day. There will be a couple of contacts each day with the ground station, with a minimum ground contact time of 25 minutes per day.

³ The observed SC attitude often shows a relative pointing error better than 1" RMS




Figure 7: CHEOPS spacecraft (payload and platform) courtesy of Airbus Defence and Space.




Figure 8: Two different views of the CHEOPS telescope.

CHEOPS in orbit and the observable sky

CHEOPS is in a Sun-synchronous Low Earth Orbit (LEO) at an altitude of 700 km. The corresponding orbital period is ~100 minutes (98.7 minutes). The satellite follows as closely as possible the day-night terminator (the Local Time of the Ascending Node of the orbit is 6 a.m.) and observes the target above the night side of the Earth (i.e., the Sun is always at the back of the satellite), therefore minimising the Earth straylight contamination. This orbit also minimises Earth occultations and, consequently, it provides a thermally stable environment. To allow the stringent thermal control of the detector, the spacecraft is 3-axis stabilised but nadir locked (see Figure 9), requiring that the payload radiators are always facing away from Earth to cold space. A small sun shield prevents illumination of these radiators by the Sun, therefore providing a thermally stable environment for the payload radiators. The fact that the spacecraft is nadir locked implies that the background stars rotate around the optical axis of the telescope, completing a circle in one orbit of the satellite. Therefore, as illustrated in Figure 9, only the target star (pointing centre) remains in the same position on the CCD over the course of an orbit, perturbed only by the pointing jitter.




Figure 9: Orientation of the satellite along one orbit. The spacecraft is nadir-locked to preserve the thermal stability. The consequence of this for CHEOPS observations is that the field in CHEOPS images (the four insets) is rotating around the target star in the centre.

To avoid the illumination of the radiators, the Sun must be outside the cone around the line of sight (LOS) of the telescope which has a half-angle of 120° . This is known as the Sun exclusion angle (Figure 10). This hard constraint in the telescope pointing limits sky coverage, making it impossible for CHEOPS to point towards the ecliptic poles.

To achieve the required photometric performance, there is an additional number of constraints on directions in which the spacecraft can be pointed. It is especially important to minimise the Earth straylight illumination. In theory, we could tolerate a minimum angle between the line-of-sight and any (visible) illuminated part of the Earth limb of 35° . This is called the Earth straylight exclusion angle (Figure 10). However, given the stray light rejection of the baffle, it is possible to maintain photometric performance at much smaller angles than this, $\sim 25^\circ$ on average. Another source of significant straylight is the Moon. An exclusion angle, the Moon exclusion angle (Figure 10) has been defined such that the bright moon must not be inside a cone with a half angle of 15° around the line-of-sight of the telescope.

The Sun, Moon and Earth stray light exclusion angles are instrument-imposed constraints. In addition, there are two other sources of interruptions to observations that are a consequence of the orbit of CHEOPS. The first is that of Earth occultations (Figure 10). For the target to be observable, its projected altitude from the surface of the Earth should be equal or higher than

100 km. The second is the South Atlantic Anomaly⁴ (SAA; Figure 10). A target is typically observable whilst flying through the SAA, however the data taken will show an increased number of glitches and noise spikes due to the impact of high energy particles. Data taken during IOC showed that images acquired during SAA passage are not amenable to photometric reduction due to the huge number of cosmic ray hits. Therefore, the images taken during SAA passage are not downlinked to ground in order to save bandwidth.

Figure 10 illustrates the constraints on sky observability described above.

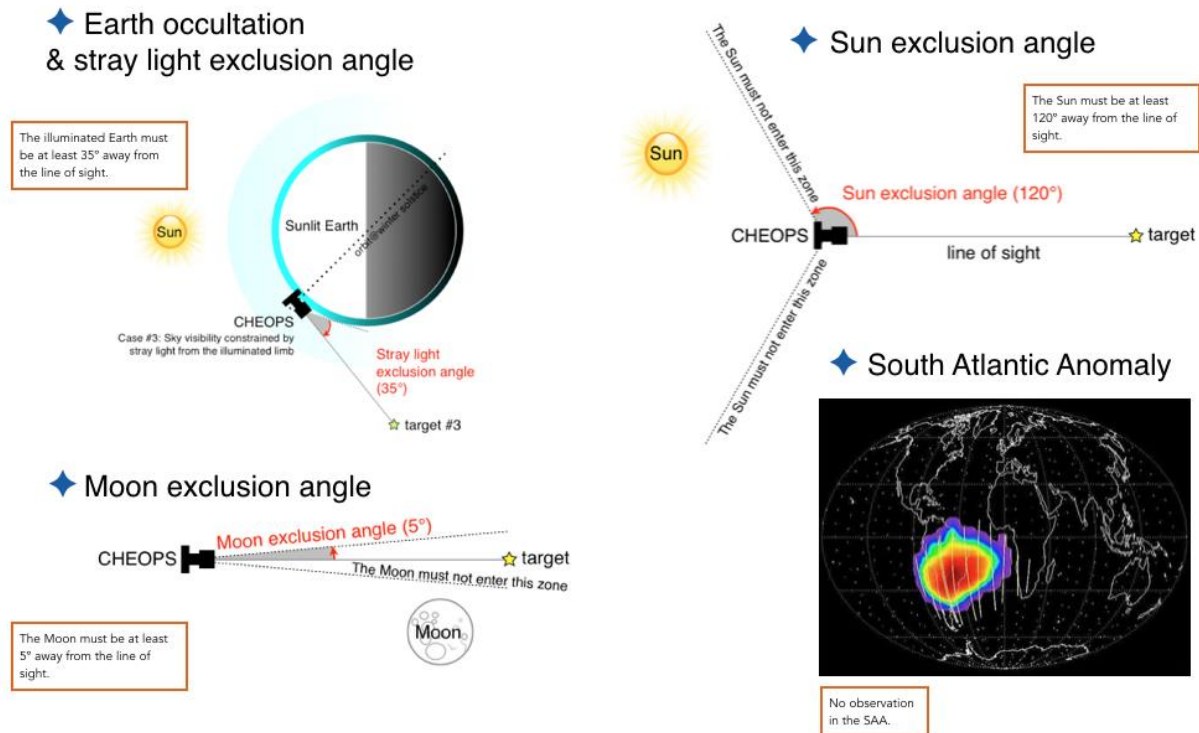


Figure 10: Schematic summary of the conditions for target observability and scientifically valid data acquisition.

The sky coverage accessible to CHEOPS can be calculated considering the constraints detailed above. Figure 11 shows a map of the sky coverage over one year. The accumulated time that can be spent pointing towards a given location in the sky (see colour code) is calculated considering only the orbits that have less than 50% of interruptions in the observations (see left panel). Calculations show that 63% of the sky is accessible for 50 cumulative days per year and

⁴ The **South Atlantic Anomaly** (SAA) is an area where the Earth's inner Van Allen radiation belt comes closest to the Earth's surface, dipping down to an altitude of 200 kilometres. This leads to an increased flux of energetic particles in this region and exposes orbiting satellites to higher-than-usual levels of radiation (from Wikipedia).

per target with time spent on target and integrating the target flux longer than 50% of the spacecraft orbit duration. The right panel shows the sky coverage over one year assuming that interruptions in the observations represent less than 20% per orbit. 22% of the whole sky is accessible for 13 cumulative days per year and per target, with the time spent on target and integrating the target flux longer than 80% of the spacecraft orbit duration.

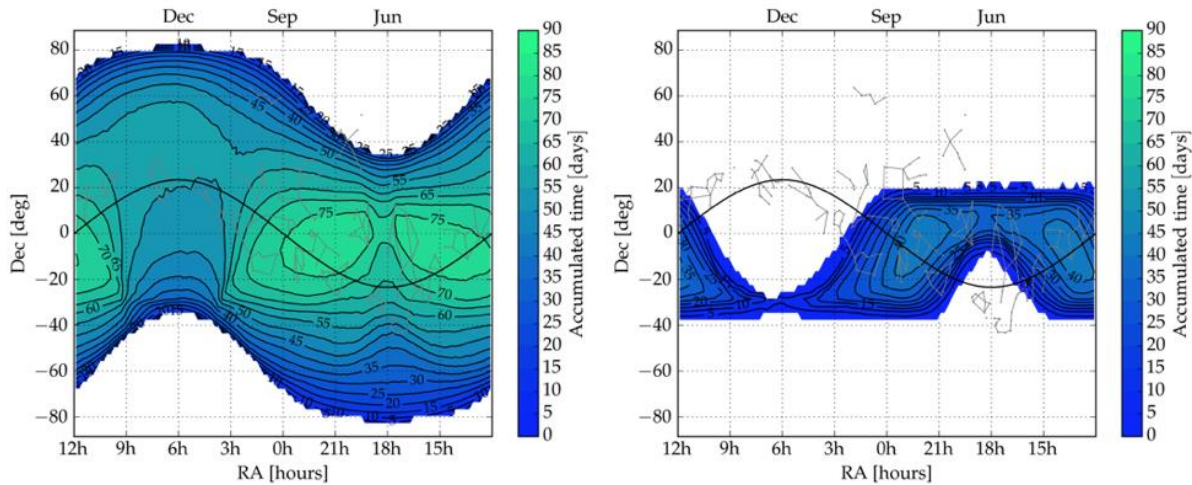


Figure 11: Left: Accessible sky in days per year for which the time spent on-target and integrating the target flux is greater than 50% of the spacecraft orbit duration. Right: same as left figure, but for which the time spent on-target and integrating the target flux is greater than 80% of the spacecraft orbit.

Optical telescope and baffle

The optical design is based on a Ritchey-Chrétien style, on-axis telescope with a focal ratio $f/5$ and an effective focal length of 1600 mm. A schematic design of the telescope is shown in Figure 12.

The CHEOPS instrument consists of 4 different units mounted on the spacecraft structure, which are:

- The baffle and cover assembly;
- The optical telescope assembly, made by the telescope itself, the focal plane and the radiators;
- The back-end electronics;
- The sensor electronics module.

The first two units are mounted on the top panel of the spacecraft, while the second two are mounted inside the structure.

The telescope has a cover designed to protect the telescope and optics from dust contamination during pre-launch activities and launch. The cover was opened during the instrument in-orbit

commissioning. Opening the cover was the only mechanical operation done in flight, as there are no other moving components in the instrument (in particular, there is no mechanical shutter). The optical system is formed by two mirrors followed by the back-end optics. The field of view has a diameter of 0.32° . The entrance pupil of the system is located at the primary mirror and has a physical diameter of 320 mm. The telescope effective diameter is 300 mm. The telescope provides a defocused image of the target star with a point spread function (PSF) of a radius of ~ 16.5 pixels (the radius of the PSF is calculated as the radius of a circle that contains 90% of the flux). A defocused PSF is used to minimise the impact of combination of the pointing jitter and the flat-field uncertainties (that should be lower than 0.1%). Figure 13 is an image of the PSF obtained during in-orbit commissioning.




Figure 12: Vertical cut of the CHEOPS telescope showing its main components.




Figure 13: Image of the CHEOPS point spread function (PSF) measured during in orbit commissioning. Left: flux distribution, x and y axis in pixels. Right: PSF profiles along the y axis for different horizontal cuts (colour curves). The y values chosen for the figure contain the pixels with the highest flux of the PSF.

To guarantee the stability of the PSF, the thermal excursions in the telescope are constrained to $\pm 1^\circ\text{C}$. The PSF is triangular due to the primary mirror mounting at 3 interface points. When cooling the primary mirror from room temperature to its operational temperature of -10°C , the mounting distorts the surface of the mirror. Due to the defocused image, this effect is very prominent.

The CHEOPS bandpass covers the visible–near-infrared range of 330–1100 nm. At the blue end, the bandpass is limited by the transmission of the optical train. In the red, it is limited by the CCD quantum efficiency. The telescope throughput over this bandpass is the product of the two terms, the optical transmission and quantum efficiency. Figure 14 shows, in the left panel, the expected quantum efficiency and optical throughput as a function of wavelength, and on the right, the global throughput (i.e. fraction of photons effectively converted into photoelectrons, integrated over the CHEOPS bandpass) as a function of the spectral type of the star.

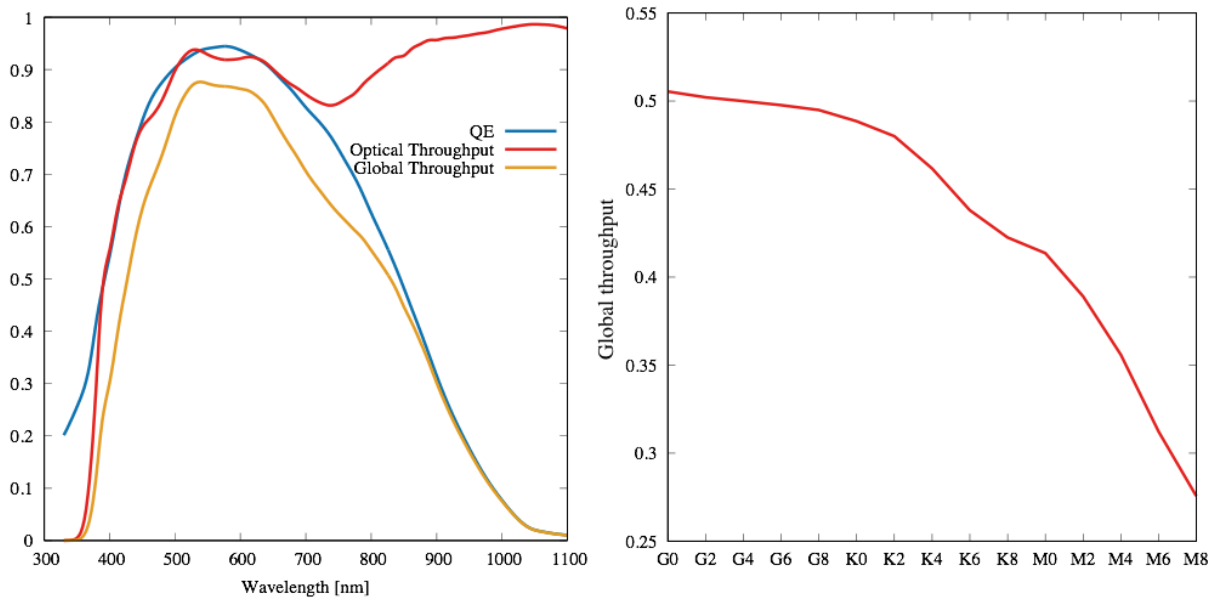


Figure 14: (Left) Quantum efficiency of the detector (blue) and optical throughput (red) as a function of wavelength. The global throughput as shown here (yellow) is just the product of the QE and Optical Throughput. (Right) Global throughput as a function of the spectral type of the star.

The telescope is designed to reject Earth reflected light which can be a strong limitation on photometry performed from a low Earth orbit. At high-incidence angles, CHEOPS can reject or attenuate Earth stray light by factors of 10^{-12} . This performance is achieved thanks to a dedicated field stop, and a baffling system comprising the baffle assembly and the telescope itself.

Focal Plane Module

The detector is a single frame-transfer back-side illuminated E2V AIMO CCD47-20. The image section of the CCD has an area of 1024×1024 pixels, while a full frame image including the covered margins has 1076 columns and 1033 rows. The camera system, which consists of the Focal Plane Module (FPM) and the Sensor Electronics Module (SEM), is designed and built by DLR Berlin, while the Back-End Electronics (BEE) is realised by IWF Graz.

The camera is controlled by the instrument main computer, named back-end electronics, using a SpaceWire link for high data transmission capability. In addition to the data link and software control, the back-end electronics delivers stabilised voltage lines to the sensor electronics module, which in turn produces ultra-stable voltages to ensure the bias voltages stabilities for the camera. This is one of the key design drivers to meet the very low noise and high gain stability requirements of the instrument.

Due to the limited telemetry budget of CHEOPS, it is not possible to continuously download full frame images. Instead, window images of 200×200 pixels with the target star in the centre, together with the corresponding covered margins are extracted from the full frame image. Note

that to further save bandwidth, the on-board software crops the corners of the window images so that circular window images (of diameter 200 pixels) are sent down to ground. Figure 15 shows a schematic of the CCD, full frame image and window image.

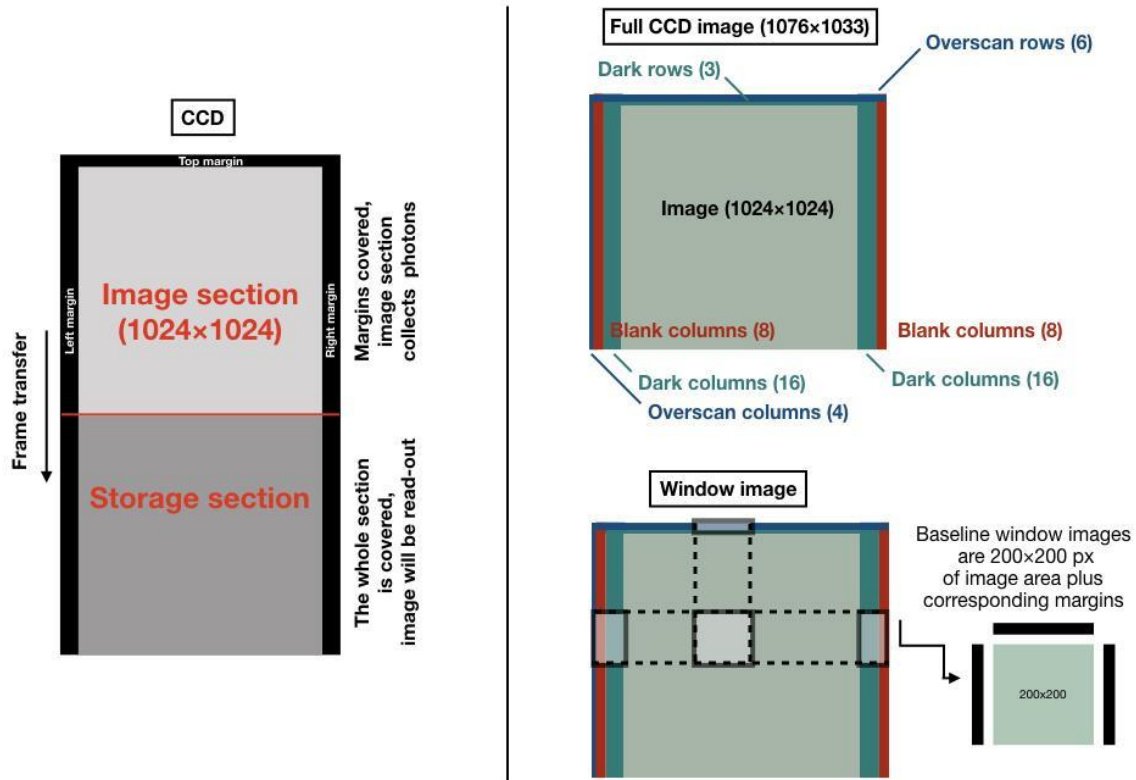


Figure 15: (Left) Schematic view of the CCD. (Right) On top, a diagram of the main elements in a full CCD image. At the bottom, a representation of a window image and its corresponding margins.

To achieve the required photometric performance, this detector is thermally stabilised to better than 2 mK (rms) at an operating temperature of -45°C . At such low temperatures, the mean dark current is negligible (less than 0.08 electrons/pixel/second at the beginning of life). Stable thermal balance is achieved by heating the instrument to precisely compensate its cooling via radiators unexposed to the Sun or the sunlit Earth.

Options for image acquisition and image processing

Operating the CHEOPS instrument consists of uplinking a set of telecommands with designated execution times through the so-called mission timeline (MTL). To command the instrument, the platform forwards the mission timeline telecommands directly to the instrument at the appropriate time without modification or addition. After the commands are executed, the platform stores and forwards the telemetry received from the instrument (science data included) to ground.

The operation of the instrument is organised according to modes. Different modes are used for instrument configuration, diagnostics, and science observations. **The SCIENCE mode is the only one relevant to CHEOPS users.** Operations in SCIENCE mode can be subdivided into several *observation activities*, each of which has an associated combination of settings for detector read-out and on-board processing of the read-out image.

The read-out modes listed in Table 1 naturally define two different kinds of products that can be downlinked to ground: *full-array images* (or *full-frame CCD images*) and *window images*. Nominal science images are circular window images of 200 pixels in diameter.

The CCD_FULL read-out mode will produce full-array images that are transmitted in their entirety (1076×1033 pixels) to the digital processing unit (DPU). This read-out mode (that works at a read-out frequency of 230 kHz) is always used at the beginning of each visit for the acquisition of the target (no images are sent to the ground during this phase) and for the calibration image (one full frame image is sent to the ground at the beginning of each visit).

The ULTRABRIGHT, BRIGHT, FAINT_FAST and FAINT read-out modes will produce window images. The read-out mode used to acquire all the images of a visit will depend on the exposure time selected (see Table 1).

Users cannot choose the read-out mode to be used during a visit; however, users can specify the exposure time of the images and the read-out mode will be automatically selected accordingly.

For all the window modes, the digital processing unit receives the CCD cropped window that was exposed to the target star (exposed window) together with the associated margins of the CCD.

Table 1 summarizes the main characteristics of each read-out mode. Note that the lowest read-out noise corresponds to exposure times longer than 2.33 seconds (read-out modes FAINT_FAST and FAINT).

Table 1: Read-out modes.

Read-out mode	Read-out frequency (kHz)	Exposure time range (seconds)	Read-out noise (ADU/px) (e^- /px)	Description
CCD_FULL	230	0.001–600	7.2 ADU/px ($\sim 15 e^-$ /px)	For full frame images. There is only one mode (this one) for reading out full frame CCD images (1024×1024). For images with exposure times lower than 5 s, the read-out is done after the exposure, and it includes the pre-cleaning of the image section. For exposure times greater than 5 s, the read-out and the exposure are done in parallel (rolling mode without pre-cleaning).
ULTRABRIGHT	230	0.001–1.05*	7.2 ADU/px ($\sim 15 e^-$ /px)	For window frame images. Only a stripe that contains the defined window plus the 9 rows of the top margin image is read-out. The read-out is done after the exposure (including pre-cleaning). <i>This means that the time between two exposures is 1.1 second (read-out time for this mode).</i>
BRIGHT	230	1.05*–2.33*	7.2 ADU/px ($\sim 15 e^-$ /px)	For window frame images. Only a stripe that contains the defined window image plus the 9 rows of the top margin is read out. Exposure and read-out in parallel (rolling mode without pre-cleaning).
FAINT_FAST	100	2.33*–60.0 (Until 30.1.23 -- 11.3*)	3.5 ADU/px ($\sim 7 e^-$ /px)	For window frame images. Only a stripe that contains the defined window image plus the 9 rows of the top margin is read out. Exposure and read-out in parallel (rolling mode without pre-cleaning).
FAINT	100	11.3*–600	3.5 ADU/px ($\sim 7 e^-$/px)	For window frame images. The full CCD is read out. The windows are cut off by SW. Exposure and read-out in parallel (rolling mode without pre-cleaning). This mode is no longer used.

CHEOPS-UGE-PSO-MAN-001

**This value corresponds to a window of 200 lines. Windows with less lines will require shorter read-out times, and windows with more lines, longer. Note that changing the size of the window frame is not an offered setting but the window size could be changed by the operators to optimize the bandwidth.*

Due to the limited data volume that can be downlinked to ground it is not always possible to download all the images obtained during each visit. Window images will be downloaded to the ground with a cadence of at least one image per minute, which can be less or equal to the acquisition cadence, which is $1/t_{\text{exp}}$ (if $t_{\text{exp}} \geq 1.05$ s, see clarification below for $t_{\text{exp}} < 1.05$ s). It is, therefore, necessary to combine the individual exposures on board to reduce the amount of data transmitted. The individual exposures are co-added at the pixel level to realise a single image. Combining (or stacking) images is the baseline working strategy of a nominal visit, although the user cannot specify how many images are stacked. As a result, the cadence of the time series of stacked images available to the observer is $1/(N_{\text{stack}} \times t_{\text{exp}})$ and is between ~ 23 s and 60 s depending on the exposure time and read-out mode. Cadence values as a function of exposure time can be found in Table 2.

To preserve some raw information about the target star, small window images (imagerettes) containing only the PSF of each image are also sent to ground. These imagerettes are also circular and typically 60^5 pixels in diameter. They are used for correction purposes in the DRP.

In non-standard observations (like monitoring and characterisation activities), the nature of the observation may require that individual images are stored on-board and then downloaded to ground. This mode is more demanding in terms of downlink budget and cannot be used as baseline but could nevertheless be used punctually for calibrations and for some monitoring and characterisation observations. Its main advantage is to transfer the whole time-resolved data set to the ground, thus preserving all information.

The nominal science activity is a cyclic process that produces data following these steps:

- Collect photons in the image section of the CCD for a defined exposure time.
- Transfer and read image with the appropriate read-out mode.
- Apply on-board image processing method (calculate the centroid, extract the imagerette in case of stacking of images, stack images, etc.).
- Pack data together with housekeeping parameters and compress.
- Send packed data to platform on-board computer.
- Repeat until the visit duration time is over.

With every contact with the ground station, the acquired data is transmitted to ground.

⁵ During IOC and early mission phases the imagerette diameter was changed from 35 to 50 and finally to 60 pixels. This was done to make use of the available bandwidth.

Baseline for the on-board data treatment and data flow

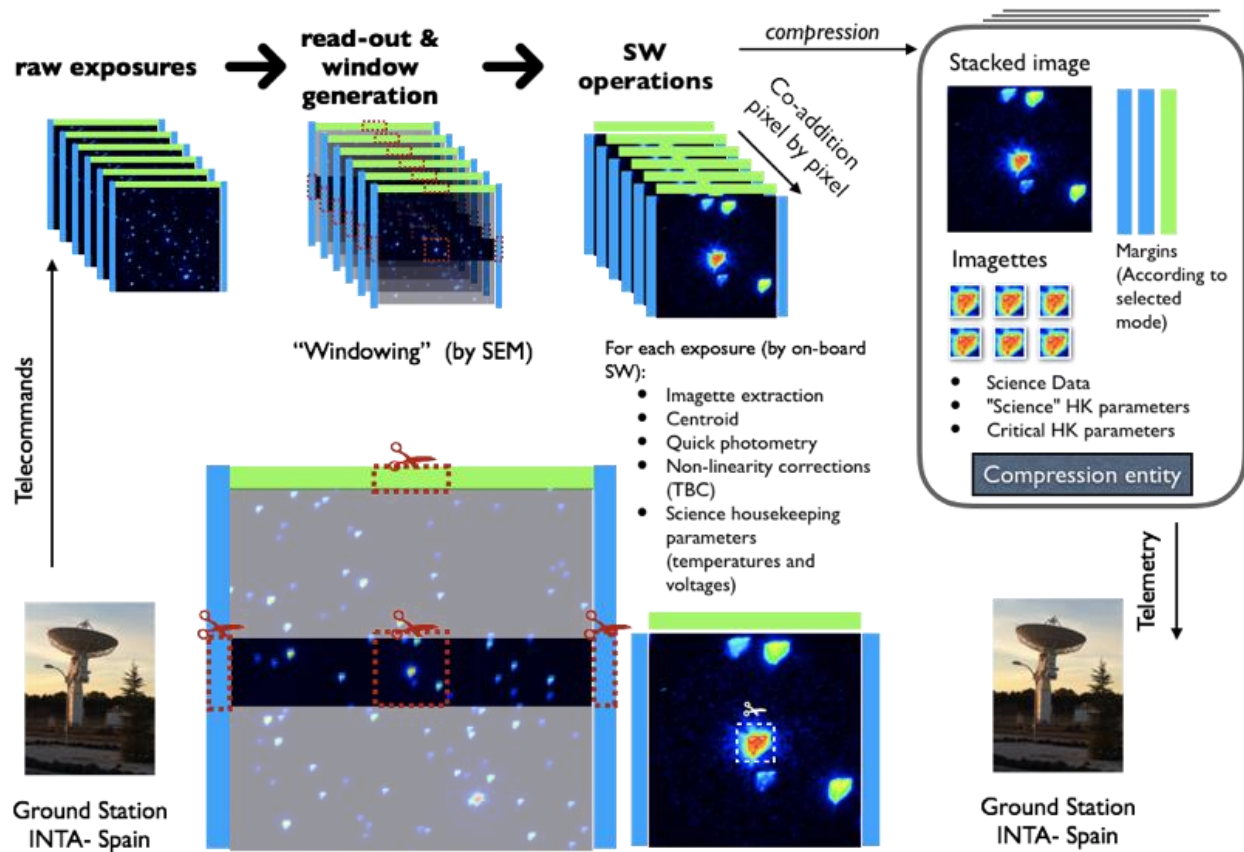


Figure 16: Sketch of a nominal science activity. Non-linearity corrections are not applied on-board, rather by the DRP (see Section 3 for details).

Figure 16 shows a sketch of the science sequence during one visit. As shown in Figure 15, each window image comes with the corresponding CCD margins. Prior to co-addition of the images, the following operations must be performed on the image:

- Centroid determination.
- Discard image (e.g. if it is taken during Earth occultation or SAA crossing).
- Quick photometry (flux measured inside the aperture to calculate the centroid).
- Extract imagette.

There are a number of different ways in which the data in the margins can be processed and stored. The default scenario, which preserves important information and reduces the amount of data as much as possible, consists of the following steps:

- The CCD margins are stacked in the same way as the image section.
- The stacked dark columns at both sides of the window are collapsed to three columns (average, median and standard deviation columns pixel-to-pixel).

- From the stacked overscan columns, the mean, median, mean absolute deviation and standard deviation are calculated.
- All blank columns are discarded.
- The three stacked dark rows are preserved.
- The six stacked overscan rows are preserved.

In addition, the following information extracted from each individual image before stacking will be downloaded:

- Centroid value.
- Quick photometry (median flux measured inside 3 apertures).
- Background estimation.
- Housekeeping parameters: temperatures and voltages relevant for image processing.

Finally, a science data packet (called a compression entity) containing

- the stacked image;
- the CCD margins information;
- the imagerettes;
- the housekeeping data associated with the imagerettes (temperatures and voltages);
- the science data (centroid, flux measured in the aperture, etc.);

is compressed and stored for download.

To further reduce the size of the compression entity, the corners of the images and imagerettes will be cropped after stacking to produce circular images (instead of squares). By only preserving the useful information, the amount of data to downlink is reduced by ~20% (see Figure 17).




Figure 17: Example of circular image extracted from a square image.

Short exposure times (less than about 3 s) result in the generation of a large number of images per minute. Although the individual images will be stacked in one single image to reduce the data volume, the individual imagerettes are, in these cases, the major contributors to the data volume. To be able to keep the data rate of one image per minute, individual imagerettes also need to be stacked.

As an example, in the case of $t_{\text{exp}} = 1.05$ s, 44 images will be stacked together to satisfy a cadence of 1 image per ~ 46 s, and imagettes will be also stacked in sets of 4 imagettes each (see Table 2), i.e. there will be 11 imagettes during the same time. See next section for more details.

Things to keep in mind when preparing your OR

There are three caveats users should have in mind when preparing their ORs.

1.1.1.8 Low duty cycle for ULTRABRIGHT mode

The ULTRABRIGHT read out mode (used for exposure times shorter than 1.05 seconds) has a low duty cycle (<50%), in the sense that between the acquisition of one image and the next there is 1.1 second used to read out the CCD.

As an example, let us assume that the exposure time is 0.4 seconds. This means that the repetition period (or time elapsed between the acquisition of one image and the next) is $0.4 + 1.1 = 1.5$ s. Therefore, one stacked image resulting from the co-addition of 30 images will be obtained every 30×1.5 s = 45 s instead of every 30×0.4 s = 12 s if the detector read-out were performed in parallel with the integration. The duty cycle for this exposure time is thus $12/45 \approx 27\%$. In conclusion, the observer is free to select the exposure time that best suits their science case but should keep in mind that exposure times of 1.04 s produce, roughly, an image every two seconds, while if the exposure time is set to 1.05 s, an image will be produced roughly every second. Select the exposure time wisely! (Refer to Table 2 for this.)

1.1.1.9 Close companions

Be aware of your target stellar field! If the target lies close to one or more other stars of similar magnitude, the photometric performance might be degraded. Even if the observation is technically feasible — which is not warranted (see below) — the photometric extraction might not be as good as in a single star case. Therefore, either avoid stars with close companions of similar magnitudes (i.e., overlapping or nearly overlapping PSFs) or be prepared for some light curve detrending (see Figure 18 for an example).




Figure 18: Example of resulting CHEOPS photometric light curve (right panel) when the target, a V=9 star, has a nearby companion of similar brightness (here at ~ 20 arcsec), as can be seen in the window image (left). The image is simulated with CHEOPSim and the light curve extraction is performed with the DRP. The visible periodic pattern in and around a transit is due to the field rotation, which makes the companion perform a full revolution around the target every satellite orbit.

The Science Operations Centre will check that the observation is feasible from a technical point of view, meaning that the target can be acquired by the instrument and that it can be kept in the centre of the image throughout the visit. The acquisition algorithm is relatively robust, but the centroiding algorithm, which basically calculates the barycentre of light within the aperture, can be fooled by close binaries with similar magnitudes. Such situations are more likely to occur for targets with faint magnitudes and could be problematic for targets fainter than the nominal magnitude range of CHEOPS ($6 \leq V \leq 12$). This may lead to a degradation of the quality of the photometry below the science requirements.

1.1.1.10 Number of images stacked on-board

The number of images and imagerettes stacked on board is a function of the exposure time. For all the cases, the cadence of the stacked image will be better than 1 stacked image every 60 seconds. The table below (see Table 2) can be used as a guide to prepare your OR in a smart way. It provides the number of images (image stacking order) and imagerettes (imagerette stacking order) that will be stacked as a function of the exposure time.

For example, if the exposure time is 1.5 s, 30 individual images and 3 individual imagerettes will be stacked on board. Every 45 seconds, 1 stacked image and 10 stacked imagerettes will be sent to ground. Keep in mind that you can only select the exposure time (the image and imagerette stacking order were calculated so as not to exceed the telemetry budget). A stacking order of 1 means that there is no stacking (stacked image[tte] = individual image[tte]). An imagerette stacking order of 0 means there is no imagerette sent to ground (not needed because images are not stacked).

Table 2: Image and imagerette stacking order, image cadence and duty cycle as a function of the exposure time. An image cadence of f means that one image is recorded every f seconds. In ULTRABRIGHT read-out mode (shaded rows), the detector has to be read-out sequentially and not in parallel to the exposition, introducing a significant decrease of the duty cycle, calculated as $d = t_{\text{exp}} / (t_{\text{exp}} + 1.1 \text{ s})$. See Table 1 for details. Mind the gap in duty cycle between exposure times of 1 s and 1.05 s!

Exposure time (s)	Image stacking order	Imagerette stacking order	Stacked image cadence (s)	Duty cycle (%)
$t_{\text{exp}} < 0.1$	40	4	$f < 48$	$d < 8.3$
$0.1 \leq t_{\text{exp}} < 0.15$	39	3	$46.8 \leq f < 48.8$	$8.3 \leq d < 12$
$0.15 \leq t_{\text{exp}} < 0.2$	36	3	$45 \leq f < 46.8$	$12 \leq d < 15.4$
$0.2 \leq t_{\text{exp}} < 0.4$	33	3	$42.9 \leq f < 49.5$	$15.4 \leq d < 26.7$
$0.4 \leq t_{\text{exp}} < 0.5$	30	3	$45 \leq f < 48$	$26.7 \leq d < 31.3$
$0.5 \leq t_{\text{exp}} < 0.55$	28	2	$44.8 \leq f < 46.2$	$31.3 \leq d < 33.3$
$0.55 \leq t_{\text{exp}} < 0.65$	26	2	$42.9 \leq f < 45.5$	$33.3 \leq d < 37.1$
$0.65 \leq t_{\text{exp}} < 0.85$	24	2	$42 \leq f < 46.8$	$37.1 \leq d < 43.6$

$0.85 \leq t_{\text{exp}} < 1.05$	22	2	$42.9 \leq f < 47.3$	$43.6 \leq d < 48.8$
$1.05 \leq t_{\text{exp}} < 1.1$	44	4	$46.2 \leq f < 48.4$	$d = 100$
$1.1 \leq t_{\text{exp}} < 1.2$	40	4	$44 \leq f < 48$	$d = 100$
$1.2 \leq t_{\text{exp}} < 1.25$	39	3	$46.8 \leq f < 48.8$	$d = 100$
$1.25 \leq t_{\text{exp}} < 1.3$	36	3	$45 \leq f < 46.8$	$d = 100$
$1.3 \leq t_{\text{exp}} < 1.5$	33	3	$42.9 \leq f < 49.5$	$d = 100$
$1.5 \leq t_{\text{exp}} < 1.6$	30	3	$45 \leq f < 48$	$d = 100$
$1.6 \leq t_{\text{exp}} < 1.65$	28	2	$44.8 \leq f < 46.2$	$d = 100$
$1.65 \leq t_{\text{exp}} < 1.75$	26	2	$42.9 \leq f < 45.5$	$d = 100$
$1.75 \leq t_{\text{exp}} < 1.95$	24	2	$42 \leq f < 46.8$	$d = 100$
$1.95 \leq t_{\text{exp}} < 2.15$	22	2	$42.9 \leq f < 47.3$	$d = 100$
$2.15 \leq t_{\text{exp}} < 2.4$	20	2	$43 \leq f < 48$	$d = 100$
$2.4 \leq t_{\text{exp}} < 2.7$	18	2	$43.2 \leq f < 48.6$	$d = 100$
$2.7 \leq t_{\text{exp}} < 2.8$	16	2	$43.2 \leq f < 44.8$	$d = 100$
$2.8 \leq t_{\text{exp}} < 2.9$	15	1	$42 \leq f < 43.5$	$d = 100$
$2.9 \leq t_{\text{exp}} < 3.05$	14	1	$40.6 \leq f < 42.7$	$d = 100$
$3.05 \leq t_{\text{exp}} \leq 3.2$	13	1	$39.7 \leq f < 41.6$	$d = 100$
$3.2 \leq t_{\text{exp}} < 3.4$	12	1	$38.4 \leq f < 40.8$	$d = 100$
$3.4 \leq t_{\text{exp}} < 3.65$	11	1	$37.4 \leq f < 40.2$	$d = 100$
$3.65 \leq t_{\text{exp}} < 3.9$	10	1	$36.5 \leq f < 39$	$d = 100$
$3.9 \leq t_{\text{exp}} < 4.25$	9	1	$35.1 \leq f < 38.3$	$d = 100$
$4.25 \leq t_{\text{exp}} < 4.7$	8	1	$34 \leq f < 37.6$	$d = 100$
$4.7 \leq t_{\text{exp}} < 5.25$	7	1	$32.9 \leq f < 36.8$	$d = 100$
$5.25 \leq t_{\text{exp}} < 6.05$	6	1	$31.5 \leq f < 36.3$	$d = 100$
$6.05 \leq t_{\text{exp}} < 7.25$	5	1	$30.3 \leq f < 36.3$	$d = 100$
$7.25 \leq t_{\text{exp}} < 9.2$	4	1	$29 \leq f < 36.8$	$d = 100$
$9.2 \leq t_{\text{exp}} < 12.5$	3	1	$27.6 \leq f < 37.5$	$d = 100$
$12.5 \leq t_{\text{exp}} < 22.65$	2	1	$25 \leq f < 45.3$	$d = 100$
$22.65 \leq t_{\text{exp}} \leq 60$	1	0	$22.65 \leq f \leq 60$	$d = 100$

Photometric performance

Noise budget

All noise sources considered below are assumed independent and can thus be added quadratically. A given noise source i is called N_i . The total noise N expected after uninterrupted integration time t is:

$$N = \sqrt{\sum N_i^2}$$

In case different noise sources are obviously correlated (for instance, the noise resulting from the combined effect of jitter and flat field or from quantum efficiency [QE] and gain sensitivity to temperature), these noises are combined together into one single component N_i , which is then added in quadrature to the other noises.

The estimation of each noise contribution is achieved with analytical formula, numerical simulations or during the calibration process. When possible, analytical estimations have been validated with more realistic models via simulations performed with CHEOPSim⁶ or other numerical models.

Here, we use the PSF obtained during the calibration campaign of the instrument (February–March 2018). The radius of the photometric aperture is assumed to be 30 pixels. We use pointing jitter time series typical of what the spacecraft platform will experience in space.

The global throughput (see section 0) is calculated for different stellar spectral types as shown in Figure 14 and used in the noise budget calculation.

1.1.1.11 Noise sources

Noise sources can be separated into two major categories depending on their nature – astrophysical or instrumental – with the total noise budget the combination of all the contributors.

The contributions we consider here are:

- Astrophysical noise
 - o Photon noise contribution from the stellar target,
 - o Zodiacal light photon noise,
 - o Noise introduced by cosmic ray hits,
 - o Stray light parasitic illumination (Earth, Moon, stars; here, we only consider the Earth contribution).
- Instrumental noise

⁶ CHEOPSim (Futyan et al. 2020, A&A 635, A23) is the CHEOPS simulator software developed with the purpose of testing the instrument and injecting simulated data into the ground segment software, such as the DRP. In these simulations, the PSFs of the target and field stars are projected onto a virtual detector. Several, but not all, noise sources are included.

- PSF in combination with the pointing jitter and the flat field uncertainties,
- Read-out noise (CCD + analogue chain random noise),
- Dark current shot noise and dark current variation with temperature,
- CCD gain variability,
- Analog electronics stability,
- Quantum Efficiency variation with temperature,
- Timing error,
- Quantisation noise of the front-end electronics.

Note that stellar noise – flicker, activity, etc. – is not included in the noise budget. This is because these noises are target-dependent and difficult to estimate a priori, on a general basis.

Figure 19 represents the expected photometric performance of CHEOPS as a function of stellar V magnitude, detailing the major contributions: instrumental noise, photon noise and noise due to stray light contamination. For more detailed information, we refer the reader to Benz et al. (2020; *Exp. Astron.*; arXiv:2009.11633).




Figure 19: Photometric noise as a function of the stellar magnitude. The solid red curve labeled “ $\Delta SL = 0$ ” (no stray light contamination) represents the noise floor of CHEOPS (calculated for an M0 star). The dashed curves are examples of noise behaviour as a function of the stray light contamination variation along the satellite’s orbit. The solid red curve labeled “ $\Delta SL = 33 \text{ ph/px/s}$ ” shows the noise behaviour when the stray light contamination dominates the total noise for all magnitudes. The green region defines the conditions that satisfy the performance science requirements. After Benz et al. (2020; *Exp. Astron.*, in press; arXiv:2009.11633).

1.1.1.12 CHEOPS expected performance

Photon noise, read-out noise and quantization noise can be considered as white noises, following

$$N_{\text{white}} = \frac{N_{\text{white}}^{\text{exp}}}{\sqrt{\text{number of exposures}}}$$

where $N_{\text{white}}^{\text{exp}}$ is the noise per exposure. On the other hand, other noise sources are considered to be systematic, and their values are considered either not to average out with time (e.g. analogue electronics stability) or, if they decrease with time, this effect was estimated through dedicated simulations (e.g. PSF + pointing jitter + flat-field uncertainties).

The input parameters to calculate the noise budget are listed in Table 3.

Table 3: Input parameters for the noise budget

Parameter	Value
PSF radius	16.5 pixels
Photometric aperture radius	25 pixels
Read-out noise (CCD & electronics)	13.8 e ⁻ /pixel for t _{exp} < 2.33 s 6.7e ⁻ /pixel for t _{exp} > 2.33 s
Quantization noise	16 bits
Pixel size	13 μm
Angular scale	1 arcsec/pixel
Pixel full well capacity	114,000 e ⁻
Image size	200×200 pixels
CCD Operating temperature	228 K
Detector temperature stability	0.002 K

Dark current at 228 K	0.022 e ⁻ /pixel/s
Flat field uncertainty	0.1%
Maximum allowed Earth straylight noise ⁷	5 ppm for bright stars 75 ppm for faint stars
Zodiacal light flux	4.02 e ⁻ /pixel/s
Detector gain sensitivity with T [ppm/mK]	1.5 ppm
Detector quantum efficiency sensitivity with T [ppm/K]	1.5 ppm
Timing error	2 ppm
Analog electronics stability	1.9 ppm
Pointing jitter	4 arcsec rms
Cosmic ray flux	1 hit in the PSF/minute (1 cosmic ray affects 4 pixels on average)

The maximum Earth straylight noise allocation corresponds to a contribution to the noise budget of ~5 ppm for bright stars with V-band magnitudes $6 \leq V < 9$ and ~75 ppm for faint stars with $9 \leq V < 12$.

The noise budget is calculated for three sizing cases that cover the science requirements on photometric precision. These three cases include a bright star $G=6$ (case A), an intermediate star $G=9$ (case B) and a faint star $G=12$ (case C). The noise is estimated for typical transit timescales. The parameters that describe each case, as well as the estimated noise contribution from the different noise sources, are shown in Table 4 for the purpose of illustration. The values listed

⁷ The different levels for bright ($V \leq 9$) and faint ($V > 9$) stars come from the different requirements in terms of achievable photometric precision.

should be taken only as a guideline to understand how the weight of the different noise components is distributed. The Exposure Time Calculator and Transit SNR predictor implement this noise budget and allow users to estimate the signal-to-noise ratio of their observations (cf. §0).

Table 4: Parameters for the various noise allocations in the noise budget.

Case	A	B	C
Star T_{eff} [K]	5,500	5,500	4,500
<i>Gaia</i> magnitude	6	9	12
Exposure time [s]	2.5	40	60

Table 5 Total noise in ppm corresponding to four different timescales (30 minutes, 1 hour, 3 hours, 6 hours) for the three size cases.

	Case A	Case B	Case C
Total noise in 30 minutes [ppm]	35	48	166
Total noise in 1 hour [ppm]	19	31	128
Total noise in 3 hours [ppm]	13	19	96
Total noise in 6 hours [ppm]	10	14	86

The values given above are approximate indicators of the expected noise in the mentioned timescales for uninterrupted and detrended light curves. We observed that for bright stars these values are very close to what is estimated from real light curves. For the faint end ($G \gtrsim 11$), the values should be regarded as optimistic (note, for example, that the noise budget does not include a model for the stellar background contamination, whose impact in the light curve uncertainties varies depending on the crowdedness of the field and the brightness of the target star).

Example data sets

1.1.1.13 Data sets from the in-orbit commissioning

A selection of data sets obtained during the CHEOPS in-orbit commissioning is described at

→ <https://www.cosmos.esa.int/web/cheops-guest-observers-programme/cheops-data>

The data can be downloaded from the CHEOPS mission archive,

→ https://cheops.unige.ch/archive_browser/

following the guidelines detailed on CHEOPS data page noted above, which also provides the individual data reduction reports generated automatically by the DRP.

1.1.1.14 Simulated data sets

Here we show three examples of CHEOPS light curves. These light curves originate from images simulated with CHEOPSim. The simulated images are as realistic as possible. The photometric extraction was performed by the DRP that will be used with real CHEOPS data. The simulations addressed the study cases described below in Figure 20 to Figure 22. All simulations show a single transit event.




Figure 20: Case 1. A bright ($V=6$) Sun-like star (G2V, 5770 K) transited by an Earth-size planet on a 50-day orbital period. The total visit duration is 20 h and the exposure time is 0.5 s. Exposures have been co-added “on-board” by stacks of 40 and the cadence is 1 min (note that these simulated data do not exactly follow the figures from Table 2, which should be used for observation preparation). The gaps are interruption of the time series by Earth occultations. There is one Earth occultation every spacecraft orbit for this target.




Figure 21: Case 2. A Sun-like star (G2V, 5770 K) of $V=9$ transited by an Earth-size planet with an orbital period of 50 days. Basically, this is the same case as in Figure 20 albeit with a fainter target. The exposure time is 10 s and on-board stacks contain 6 images.




Figure 22: Case 3. A faint ($V=12$) star (K5V, 4410 K) transited by a Neptune-size planet on a 13-day orbital period. The total visit duration is 10 h and exposure time is 60 s (exposures are thus not co-added).

The data products corresponding to these simulations are available for downloading and perusing at the following address:

→ <https://www.cosmos.esa.int/web/cheops-guest-observers-programme/cheops-data>

The format, file structures and names of these simulated data files are identical to those of real CHEOPS observations, therefore readers and potential users might be interested in their manipulation. Each case data set is provided as a compressed tar file consisting of many files. Each tar file contains all products generated by SOC for a single visit. A large fraction of these files are intermediate products and log files that are not directly relevant for the end-user. We provide below some guidelines to browse through these files. The nomenclature system of CHEOPS files is further described in Section 0.

1.1.1.15 High level products containing the information most relevant for data analysis

Data Reduction Visit Report

CH_PR*_RPT_COR_DataReduction_V0000.pdf.

This report provides an illustrated summary of the data reduction steps applied by the pipeline (see Sections 0 and **Error! Reference source not found.**).

Light curve extracted for several aperture sizes

CH_PR*_SCI_COR_Lightcurve-{DEFAULT,OPTIMAL,RINF,RSUP}_V0000.fits

These files also contain flags/events, estimates for background and contaminants, roll angle and target centroid position (at the cadence of stacked images, if applicable). The aperture information is provided in the DR Visit Report and as keywords AP_RAD1 and AP_TYPE in the FITS file.

Event flags

The 'EVENT' column in the SCI_COR_Lightcurve data structures should be decoded according to the position of decimal digits in the flag: |7|6|5|4|3|2|1|0|

- pos. 0: SAA (0: outside; 1: inside) we can also define 2: border of SAA (need definition)

- pos. 1: Temperature (0: within range; 1: outside range (soft); 2: outside range (hard);
 - pos. 2: EARTH (0: above threshold; 1: below threshold (soft=35 degrees); 2: below threshold (hard=25 degrees); 9: earth occultation)
 - pos. 3: MOON (0: above threshold; 1: below threshold (soft=5 degrees); 2: below threshold (hard=5 degrees))
 - pos. 4: SUN (0: above threshold; 1: below threshold (soft=120 degrees); 2: below threshold (hard=120 degrees))
 - pos. 5: CR pixel (0: below threshold; 1: >1% aperture CR contaminated; 2: >5% frame CR contaminated)
- For example: The event 901 means: 9=earth occultation & 0=Temp OK & 1 =inside SAA (the rest of the digits to the left of the 9 are 0 = OK)

Corrected subarray data cube

`CH_PR*_SCI_COR_SubArray_V0000.fits`

This file contains the fully corrected stacked images, the metadata (WCS, angles between LOS and Sun/Earth/Moon), the target centroid position, and the smearing information. From these files the DRP extracts the light curves to `CH_PR*_SCI_COR_Lightcurve-*_V0000.fits`.

CHEOPS reduced light curves are in FITS files called
`CH_PR*_SCI_COR_Lightcurve-*_V0000.fits`.

1.1.1.16 Mid-level products, possibly useful for deeper analysis

- Target centroid position at the cadence of individual exposures:
`CH_PR*_SCI_RAW_Centroid_V0100.fits`
- Spacecraft attitude (RA, Dec, roll) at the cadence of AOCS data:
`CH_PR*_SCI_RAW_Attitude_V0100.fits`
- Imagettes data cube at the cadence of individual exposures, or at the cadence of stacked imagettes if applicable:
`CH_PR*_SCI_RAW_Imagette_V0100.fits`
- Star catalogue of Gaia field stars:
`CH_PR*_EXT_PRE_StarCatalogue-OBSID*_V0000.fits`
- Raw FullFrame array with side/top margins and metadata:
`CH_PR*_SCI_RAW_FullArray_V0100.fits`
- Information based on the restituted/measured orbit (angles, latitude/longitude, Earth occultation and SAA flags):
`CH_PR*_AUX_RES_VisitConstraints_V0100.fits`
- Quick Look reports (one per pass):
`CH_PR*_RPT_RAW_QuickLook_V0100.pdf`
- Calibrated subarray data cube:
`CH_PR*_SCI_CAL_SubArray_V0000.fits`

1.1.1.17 Low level products, mostly relevant to the SOC and the Instrument Team:

- log files: CH_*.log
- Raw subarray data cube: CH_PR*_SCI_RAW_SubArray_V0100.fits
- Raw products per pass: CH_PR*_TG*_PS*.fits
- Pre-processing visit report: CH_PR*_RPT_RAW_Visit_V0100.pdf
- House Keepings: CH_PR*_SCI_RAW_Hk*_V0100.fits
- Event report: CH_PR*_SCI_RAW_EventReport_V0100.fits
- Visit combination summary: CH_PR*_PIP_REP_Text-visit_combination_V0100.fits

1.1.1.18 Differences with real data

While these simulated data sets are largely identical to expected real data from CHEOPS, a number of differences exist. They are listed below for the sake of completion but are not foreseen to impact most possible usage of these simulated data sets.

1. Roll angle is not (yet) provided in CH_PR*_SCI_COR_Lightcurve-*_V0000.fits. Instead, the Roll angle may be retrieved at cadence of individual exposures from CH_PR*_SCI_RAW_Attitude_V0100.fits.
2. Target centroid position is provided in CH_PR*_SCI_COR_Lightcurve-OPTIMAL_V0000.fits but in the wrong column (Location_X/Y instead of Centroid_X/Y)
3. Target centroid position is currently provided relative to the bottom left corner of the window frame. It will in the future be relative to the bottom left corner of the CCD full frame. The pixel coordinates of the bottom left corner of the window frame relative to the bottom left corner of the CCD is provided in the header of the first extension of SCI_COR_SubArray as X/Y_WINOFF.
4. Smearing information is empty in CH_PR*_SCI_COR_SubArray_V0000.fits.
5. The event flagging by the DRP is still based on the angles to SUN/MOON/EARTH from the AUX_RES_VisitsConstraints, which is sampled at a 5-minute cadence. With this coarse cadence, some frames may be flagged inappropriately. For Straylight event flagging (EARTH=1 or 2), the DRP uses angles between line-of-sight and Earth to set the flags instead of using the predicted straylight levels and applicable thresholds. Some frames may be flagged inappropriately.
6. The stacking orders and cadences of simulated data differs a bit from the up-to-date figures from Table 2 (which should be used for the preparation of observations).

Calibration, monitoring & characterisation of the instrument

1.1.1.19 Ground-based calibration of the instrument

The CHEOPS instrument was calibrated on ground. This process provided reference calibration products for the bias level, read-out noise, system gain, detector pixel non-linearity and full well capacity, dark current, global throughput and flat field.

It is important to note that the majority of these on-ground activities cannot be reproduced in space.

The key results of the calibration campaign have been reproduced in this document. In particular, the measured point spread function (PSF) of the CHEOPS instrument is shown in Figure 13 and it shows the PSF as measured in space.

The detector flat-field is reproduced from Figure 23 to Figure 25, seen in filters of various passbands centred at different wavelengths. Measuring the flat-field in many different passbands, narrow and broad, allows a detailed characterisation of the flat-field structure. As can be seen in Figure 24, the bluest passbands showcase the “blue diamond” effect (also well visible in the broader *U* and *B* filters in Figure 23); this effect is due to surface imperfections on the silicon chip introduced during the manufacturing process. The reddest passbands, on the other hand, show the fringing effect, which is the result of interference between the incident light and the light internally reflected at the interfaces between the thin layers of the CCD chip. Both effects can lead to photometric variations when the target moves across the detector as a result of jitter motion. While the broad passband of CHEOPS allows to average out these effects for stellar targets with a broad spectral continuum, their respective influence is different for targets with different spectral types. This is why the DRP will tailor the flat-field correction to match the spectral energy distribution of the target, which is a user-provided information. For a target with a given spectral type, the flat-field will be a composition of the monochromatic flat-fields⁸ from Figure 24, weighted by a spectral template of the target source. The CHEOPS flat-field composed over the *V* band can be seen in Figure 25.

⁸ The monochromatic flat-fields are not used beyond ~800 nm due to fringing. Instead, we make use of the flat-fields measured through the broad-band *R* and *I* filters from Figure 23.




Figure 23: Broad-band flat-fields measured on the CCD in nominal conditions (detector at -40°C and front-end electronics at -10°C) with the U, B, V, R and I Johnson filters. The precision and scale are the same as in Figure 24. (Extracted from the CHEOPS Payload calibration report.)

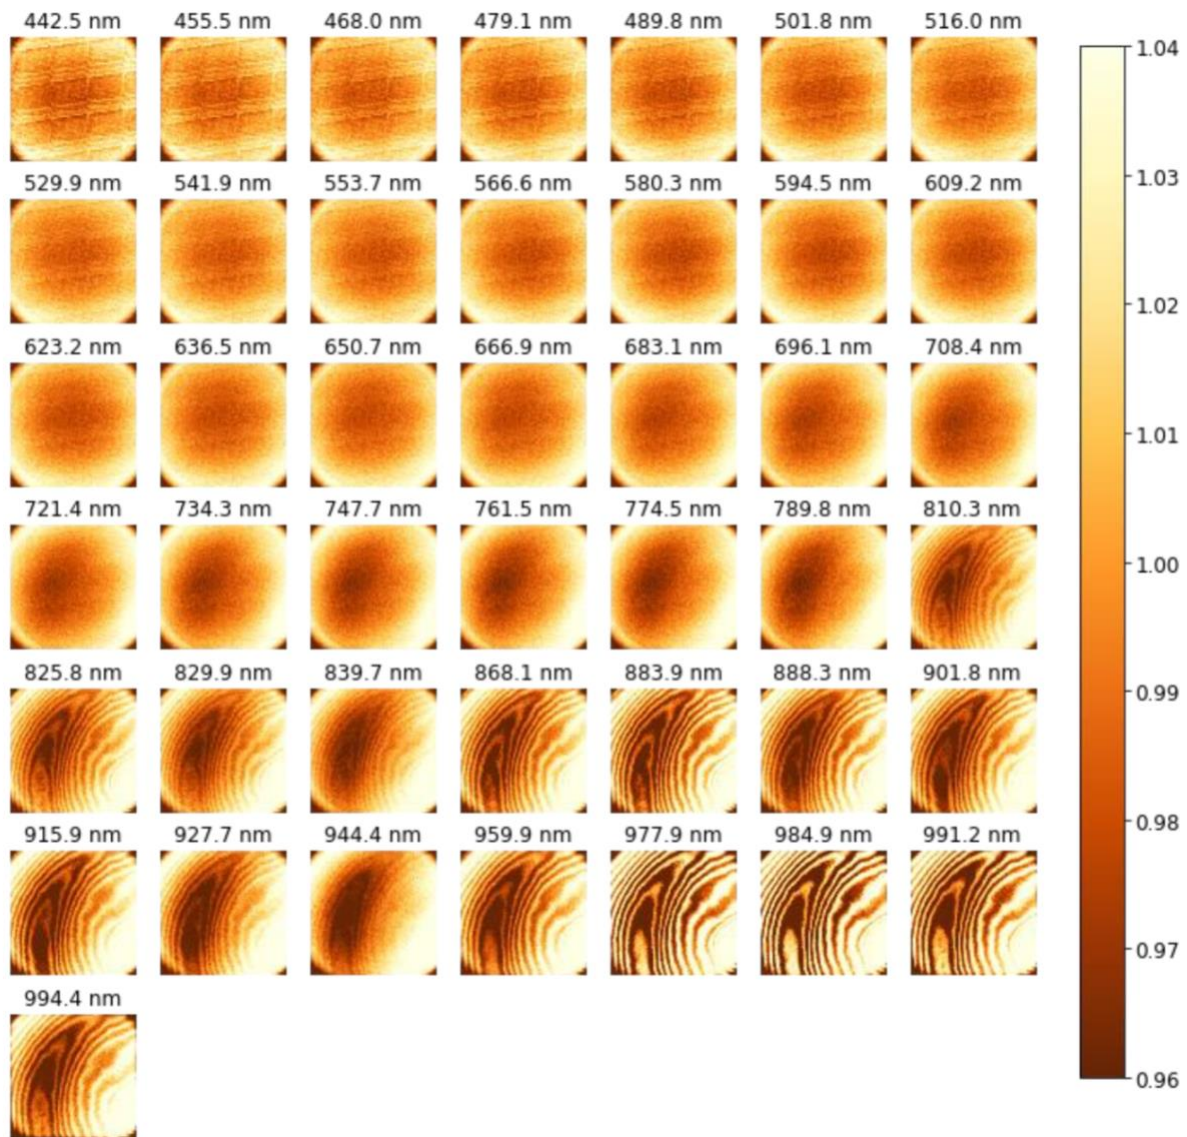


Figure 24: Monochromatic flat fields measured on the CCD in nominal conditions (detector at -40°C and front-end electronics at -10°C). The bandwidth of each measurement is ~ 30 nm. The precision on each pixel is lower than 0.1% with an average precision of 0.06%. The colour bar provides a linear scale fixed between 0.96 and 1.04 to show the structure of all flat-fields. About 90% of the values are within that range. (Extracted from the CHEOPS Payload calibration report.)

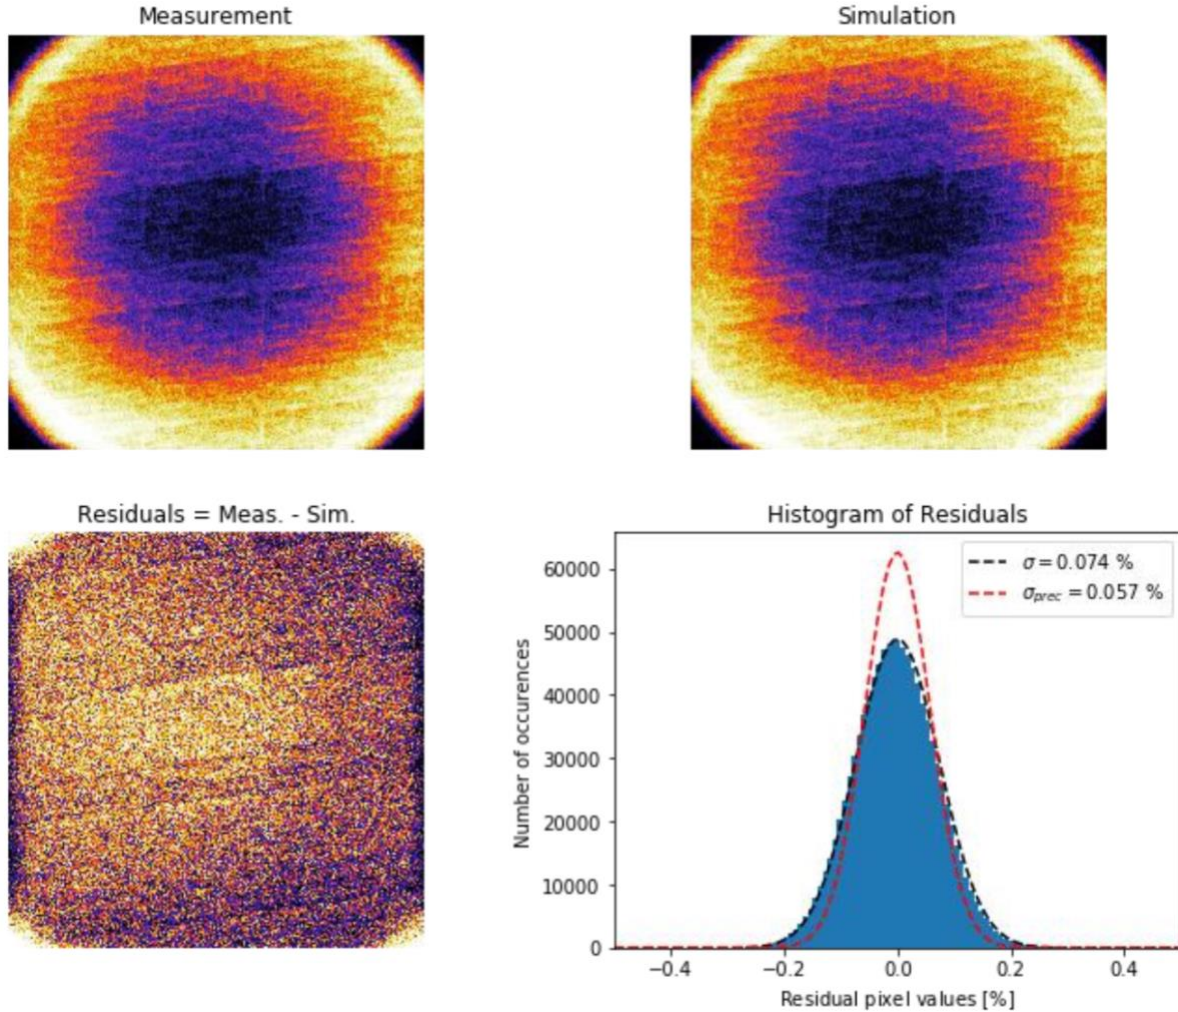


Figure 25: Result of the flat-field combination for the V filter. The upper-left image shows the flat field measured with the Johnson V filter. The upper-right image shows a flat-field synthesised as the average of the monochromatic measurements from Figure 24, weighted by the transmission spectrum of the V filter. The lower-left image is the residuals of the top-panel images. The lower-right plot shows the histogram of the residuals (in blue), with a standard deviation of 0.074%. The two dashed lines are plotted for visualisation purposes: the black one is a Gaussian distribution with the same dispersion as the residuals and the red one represents the distribution that could be reached theoretically, if we were limited only by the flat-field measurement noise (0.057%). (Extracted from the CHEOPS Payload calibration report.)

The photometric precision was measured on the calibration bench and was also ultimately limited by the calibration bench (see the Payload calibration report). Despite this, the RMS of photometric time series measured on the calibration bench reached levels on par with the science requirements (see Figure 26).




Figure 26: Noise (rms) of the residuals from a photometric time series recorded on the calibration bench, as function of the time series binning (blue curve). The sky-blue shaded area represents the $\pm 1\sigma$ uncertainty. The black dashed line indicates the slope of a white-noise dominated time series. At a time bin of 6 hours, the noise level has reached a threshold of ~ 25 ppm, likely set by the set-up of the calibration bench and not the instrument level.

1.1.1.20 In-orbit monitoring and characterisation programmes

CHEOPS has neither an on-board calibration lamp (to measure the flat-field) nor a shutter (to measure the dark current) and so cannot be calibrated in-orbit. Instead, during operations, we will monitor the instrument performances and characterise the instrumental systematic effects to the best precision possible through a series of dedicated programmes, called Monitoring & Characterisation (M&C) observations, and make use of the nominal science visits themselves. The M&C programme will undertake the following activities:

- “Pseudo” flat-fielding on a crowded field;
- Mapping of the detector background (including dark current and hot pixels) on void sky regions;
- Mapping bad or dead pixels and the level of stray light from the Earth on the illuminated Earth limb;
- Monitoring the PSF and evolution of the flux sensitivity across the detector;

- Measuring the evolution of the photometric precision with some reference transit observations;
- Mapping the full well capacity and gain of the detector pixels.

During operations, analyses of the M&C data performed by the Project Science Office will be regularly posted online as Instrument Science Reports.

Overview of the science operations

Mission planning

The CHEOPS Consortium is responsible for Mission Planning activities. The Science Operations Centre (SOC) in Geneva automatically generates optimised long-term and weekly CHEOPS observing schedules. The optimisation process is based on a so-called “merit function” used to guide the convergence of a genetic algorithm. This merit function is the weighted sum of elements such as the scientific priority (see below) of an OR, its effective duration (real on-source time, excluding interruptions due to SAA crossing or Earth constraints/occultations), or the overall balance between GTO and GO programmes of the generated timelines.

The generated mission timelines are uploaded once per week to the spacecraft. Data downlink is done whenever a ground station contact is available, typically twice per day. Each pass is processed with Quick-Look and Data Reduction is started for each visit when all related science and housekeeping data has been received.

Scientific priorities of observations

CHEOPS observations are assigned scientific priorities for scheduling purposes. These priorities are used by the schedule solver algorithm to generate optimised observation sequences. There are 3 possible grades of priorities:

1. Priority-1, is assigned to high-priority or otherwise urgent observations. The relative fraction of priority-1 observations is kept below about 30% to ensure that their chances of being scheduled remains high, i.e., to avoid scheduling conflicts (concurrent observations) for high-priority observations.
2. Priority-2, is assigned to the bulk of CHEOPS observations. These observations receive a neutral weight in the schedule optimisation process.
3. Priority-3, is assigned to low-priority observations. Due to the time-critical nature of most CHEOPS observations, and to ensure a high filling factor for the observing schedule, some level of oversubscription will be required. Most of the oversubscription will be accounted for by the priority-3 observations. Consequently, over half of these priority-3 observations will not be scheduled. Note that short non-time-critical priority-3 observations would serve well as “filler” observations and have, therefore, a higher chance of being scheduled, considering their intrinsic low scientific priority.

On-ground data processing

The SOC automatically receives the science and housekeeping telemetry after they have been received by the Mission Operations Centre. Data are automatically processed and archived at the SOC. As a first step, the binary telemetry files are converted to FITS data structures for further data processing. This conversion is done by the pre-processing pipeline. Once all data for a given visit have been received and converted to FITS data structures, the DRP will automatically process the data to produce the intermediate and final data products. The DRP and its data products are described in Section 0. The configuration of the DRP is identical for visits of the GTO and the GO programmes. After the successful processing by the DRP, all output data will be made available in the CHEOPS Mission Archive.

Data archive

Once the data for a given visit have been successfully processed by the DRP, the data will be made available via the CHEOPS Mission Archive. The PI of the relevant visits will be notified by email that data can be downloaded.

The CHEOPS Mission Archive is the unique interface to access proprietary CHEOPS data from the GO Programme's AOs and DPs. The archive can be accessed via:

→ https://cheops.unige.ch/archive_browser

CHEOPS data

CHEOPS data structure

Data levels

The data received at SOC or processed at SOC belong to well-defined *data levels*. Most of the data are stored in FITS files. The data level is defined in the FITS header keyword DATA_LVL. The different data levels are defined below.

Level 0—The data received by the SOC from the Mission Operations Centre (MOC). They contain the science data, as well as the housekeeping data and auxiliary data. All Level 0 data received at SOC are directly archived and used as inputs for the data processing chains. The format of the data is either telemetry packets as downlinked from the spacecraft, XML files or plain ASCII files as produces by MOC for the auxiliary data.

Level 0.5—Data resulting from the pre-processing step at SOC. The telemetry data, science and housekeeping data, are converted into FITS files and sorted by visits to serve as input for the DRP or by visits and pass to serve as input for Quick Look. Housekeeping data are converted into physical values. The data are time tagged with UTC and MJD.

Level QL—The Quick Look processing of Level 0.5 data produces quality reports in PDF format. The data are time tagged with UTC, MJD and BJD_{TT}.

Level 1—Data resulting from the processing of Level 0.5 data by the DRP. They are calibrated and corrected science images, full-array images as well as sub-array images. They include engineering meta-data associated with the science data. The data are time tagged with UTC, MJD and BJD_{TT}.

Level 2—Data resulting from the processing of the Level 1 images by the DRP. They consist of photometric time series (light curves) and associated meta-data. These are the CHEOPS final science products, which are expected to be used as input to the science analysis by scientists.

AUX—Auxiliary data converted from the auxiliary data of Level 0 into FITS files. Sorting and merging data from several deliveries from MOC can be implemented, depending on the type of the auxiliary data.

MPS—Planning data generated by the Mission Planning System. These are the activity plan, sent to MOC as well as the planning data used by the pre-processing to assign the data, receive by MOC, to a planned visit.

REF—Reference data, obtained from on-ground calibrations and in-orbit monitoring and characterisation programmes. These data are not generated by SOC but received from the Instrument Team and the Project Science Office. They are used in various steps of the data processing at SOC.

EXT—Externally delivered data like catalogues or the contour map of the SAA.

SOC—Data generated at SOC in a manual processing step.

SIM—Simulated data created by CHEOPSim.

Types of CHEOPS files and file nomenclature

The names of CHEOPS output files follow the format:

```
CH_PRppnnnn_TGooooon_PSyyymmhhdd_TUyyyy-mm-ddThh-mm-ss_DataStructureDefinition-DataName_Vnnnn.fits,
```

where the component strings separated by an underscore (_) are described below. First, here is an example:

```
CH_PR100040_TG002004_PS19010406_TU2019-01-03T16-32-27_SCI_RAW_SubArray_V0000.fits
```

Examples of (simulated) CHEOPS data sets can be found in Section 0.

Table 3: Description of CHEOPS file names

Name	Description
CH	Abbreviation of CHEOPS.

PRppnnnn	<p>A programme ID consisting of two digit (pp) identifying the type of programme (Programme_Type), followed by a four-digit (nnnn) arbitrary number (Programme_ID), starting with 0000.</p> <p>This component is skipped in the file name if the data do not belong to a specific programme, like auxiliary data or reference data.</p>
TGooooon	<p>The visit of the target in an OR. It consists of four digits (oooo) defining the OR ID and two digits (nn) counting the visits of the target within the OR. The programme ID together with the target visit defines the visit ID.</p> <p>This component is skipped in the file name if the data do not belong to a specific programme, like auxiliary data or reference data.</p>
PSyymmhhdd	<p>The pass ID, defined by the time (year, month, day and hour) of the pass. This component is skipped in the file name if the data do not belong to a specific pass, like auxiliary data or reference data.</p>
TUyyyy-mm-ddThh-mm-ss	<p>A timestamp. It is the time of the first data in the FITS / XML file expressed in UTC. For reference data with no data time, it is the start time of validity.</p>
DataSetDefinition	<p>The name of the main data structure definition in the product. This is typically the name of the first extension but has to be defined case by case. For log files it is the program that created the log file.</p>
DataSetName	<p>The name of the data to distinguish data for the same time period in the same DataSetDefinition. This is an optional field.</p>
Vnnnn	<p>The first two digits of this version defines the archive revision while the last two digits define the processing number in case the same data have to be processed more than once.</p>

CHEOPS files can belong to the following categories.

- Files types of corrected data:
 - o SCI_COR_Lightcurve
 - o SCI_COR_SubArray
 - o RPT_COR_DataReduction
- Files types of calibrated data:

- SCI_CAL_SubArray
- Files types of raw data:
 - SCI_RAW_FullArray
 - SCI_RAW_SubArray
 - SCI_RAW_Imagette
 - SCI_RAW_Attitude
 - SCI_RAW_Centroid
 - SCI_RAW_HkCentroid
 - SCI_RAW_EventReport
 - SCI_RAW_HkAsy30759
 - SCI_RAW_HkCe-FullArray
 - SCI_RAW_HkCe-SubArray
 - SCI_RAW_HkDefault
 - SCI_RAW_HkExtended
 - SCI_RAW_Hklfsw
- Reference file types:
 - REF_APP_BadPixelMap
 - REF_APP_BiasFrame
 - REF_APP_CCDLinearisationLUT100
 - REF_APP_DarkColumns
 - REF_APP_DarkFrame
 - REF_APP_EventEnumConversion
 - REF_APP_EventParamConversion
 - REF_APP_FlatFieldFilter
 - REF_APP_FlatFieldTeff
 - REF_APP_FluxConversion
 - REF_APP_GainCorrection
 - REF_APP_HkEnumConversion
 - REF_APP_HkParamConversion
 - REF_APP_Limits
 - REF_APP_PixelScale
 - REF_APP_QE
 - REF_APP_SEDFilter
 - REF_APP_SEDTeff
 - REF_APP_Throughput
 - REF_APP_VisitConstraints
 - REF_APP_WhiteCCDLocationPSF
 - REF_APP_WhitePSF

CHEOPS file structures

Most of the CHEOPS data are provided in FITS files. The primary array is never used to store data, even if the FITS file stores an image. The Extension Name (EXTNAME) is identical to the data structure name as defined in the Data Product Definition Document (DPDD). The applicable DPDD is part of the AO documentation package available via the webpages for the Call.

The first extension in the FITS file stores the main data product of the file. Its Extension Name is repeated in the file name.

Headers, keywords and relationship to Phase 2 parameters

Several of the parameters provided in the PHT2 are copied into Header Keywords. Table 4 shows the relation between the Phase 2 parameters and the FITS Header Keywords.

Phase 2 parameter	FITS header keyword
Right Ascension	RA_TARG
Declination	DEC_TARG
Target Name	TARGNAME
Target <i>Gaia</i> magnitude	MAG_G
Target <i>Gaia</i> magnitude error	MAG_GERR
Proprietary Period First Visit	PRP_VST1
Proprietary Period Last Visit	PRP_VSTN
Effective temperature of the target	T_EFF
Exposure Time	EXPTIME

Table 4: Correspondence between FITS header keywords and Phase 2 parameters.

Errors and data quality flags

There are two columns in the light-curve table that provide information about the quality of the data:

- STATUS: A flag indicating the status of the measurements.
- EVENT: A flag indicating the possible events that might affect the measurement (without necessarily invalidating it).

Coordinate systems

1.1.1.21 Sky coordinates

In general, sky coordinates are provided in right ascension (RA) and declination (DEC) at the epoch of the actual observation. In future revisions of the archive, we plan to provide coordinates at the epoch J2000. The coordinate reference frame is ICRS. The sky coordinates of the target star and the background stars in the CHEOPS Sky Catalogue, i.e. in the data structure EXT_PRE_StarCatalogue. These are provided at the epoch of the visit (keyword OBSEPOCH). The CHEOPS Sky Catalogue used by the SOC is based on information available via the Gaia DR2.

1.1.1.22 Pixel coordinates

The following conventions apply whenever a CCD pixel is addressed by its coordinates (x,y) in a SOC data structure: The bottom left pixel of the illuminated area of the CCD (i.e. the CCD without margins) has coordinates (0,0). The coordinates of the top right pixel of the illuminated area are (1023,1023).

Pixels offsets, for example defining the location of the sub-frame window on the CCD, are always relative to pixel (0,0).

The following conventions are used whenever a location on the CCD is defined, for example the intended or the actual location of the centre of the target star: The location is specified with an accuracy of $1/100^{\text{th}}$ of a pixel, i.e. in centi-pixels. The bottom left pixel of the illuminated area (i.e. the CCD without margins) has coordinates from (0.00,0.00) to (0.99,0.99). The centre of that pixel has coordinates (0.50,0.50). The centre of the top right pixel of the illuminated areas has the coordinates (1023.50,1023.50). The centre of the illuminated CCD has the coordinates (512.00,512.00).

Time references

There are 4 different times provided with the data, depending on their *data level*.

- On-board-time (OBT)
 - o Level 0.5
- UTC
 - o Level 0.5
 - o Level 1
 - o Level 2
- MJD (TT)
 - o Level 0.5
 - o Level 1
 - o Level 2
- BJD (TT)
 - o Level 1
 - o Level 2

Data Reduction Pipeline (DRP): overview and algorithms

Pipeline processing overview

The DRP reduces the series of raw images in a given visit, the Level 0.5 data products, to produce the corresponding flux time series, ready for the scientific analyses. These light curves, called Level 2 Data Products, are then archived in the CHEOPS Mission Archive. The DRP is run automatically once triggered by the processing framework. There is no interaction with an external

agent, and there is no interactive configuration of the pipeline. The pipeline is implemented in Python 3.4.

Structure of the pipeline (module description & flowcharts)

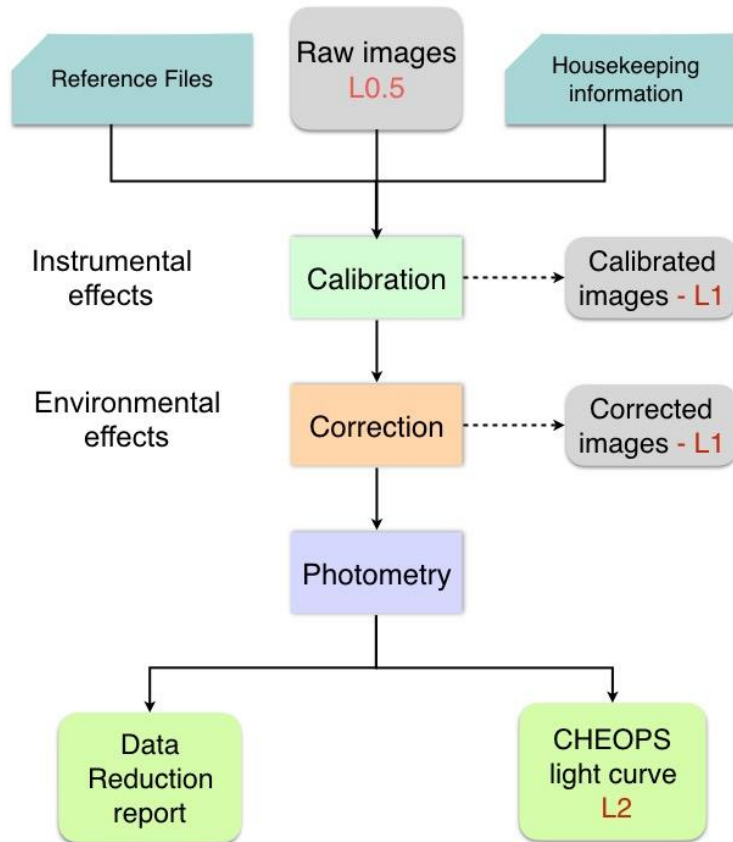


Figure 27: General flow chart diagram of the DRP.

A complete description of the DRP is provided in Hoyer et al. (2020, A&A 635, A24). Here, we provide a general overview of the different steps of the data processing chain. It follows the regular stages in photometric data reduction but is adapted to CHEOPS specificities, which are:

- the rotation of the platform inducing a rotation of the field of view around the target;
- the relatively small field of view;
- the point spread function of the instrument.

The DRP is divided into 3 main components: calibration, correction, and photometry, as shown in Figure 27 **Error! Reference source not found..** In addition to the reduced light curves and a set

of relevant time series with parameters used in the reduction/correction of the data (such as flux contamination, background, smearing, etc), the pipeline issues a report which allows the user to get direct insights on how every step of the data reduction went (see §3.2.4). The successive processing steps are described in the next sections.

In its pre-launch version, the pipeline has been established using CHEOPS simulated data and inputs from the CCD characterization done in the laboratory. Its performances have been assessed using simulated data of typical astrophysical configuration as provided by CHEOPSim. Since March 2020, it has been fine-tuned to adapt the codes to in-flight conditions using data acquired during in-orbit commissioning and the first months of scientific observations.

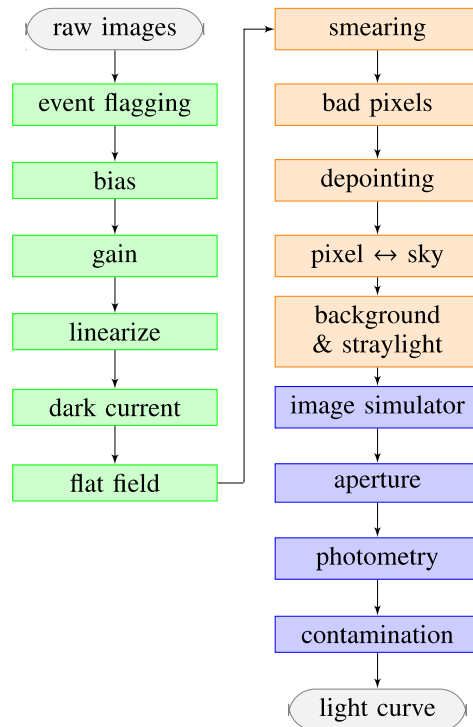


Figure 28: Flow chart diagram of the DRP with the successive reduction steps. The colour code corresponds to the three successive modules: calibration (green), correction (orange) and photometry (blue).

Data flow through the pipeline

The calibration step transforms the raw images downlinked from the instrument into calibrated photo-electrons images. It removes the bias introduced by the analogue chain, restores the ADU (Analog Digital Unit) image back to electrons, evaluates and corrects the dark current.

The correction step gathers all the modules that are correcting for environmental or instrumental effects affecting the images. Starting from the calibrated image in electrons, smearing trails

produced by the target and nearby stars are corrected, then bad pixels are identified and corrected. Finally, the background, zodiacal, and stray lights are measured and subtracted.

The last step of the reduction pipeline is the photometric extraction, during which several light curves are extracted with their associated metrics from the corrected image series. An automatic execution report is then generated, and these data are transferred to the archive.

1.1.1.23 Calibration

Inputs to this module are the Level 0.5 data products. It transforms the raw images received from the instrument into photo-electrons calibrated ones.

The sequence of processing steps in the Calibration component that produces Level 1 calibrated images and their associated uncertainties consists of the regular processing steps in photometry:

- the event flagging;
- the bias and read-out noise estimate;
- the gain and nonlinearity corrections;
- the dark current correction;
- the flat-field correction.

In addition to the 200 columns and 200 rows of physical pixels that correspond to the photo-sensitive CCD image area, special rows and columns in the margins of the image are downlinked. They consist of 4 pre-scan, 8 unused blank, and 32 dark columns plus 6 overscan and 3 dark rows.

Bias and read-out—This correction includes (i) the removal of the temporal variation due to thermal effects of both read out noise and stabilised offset voltage applied to the amplifier before analogue-to-digital conversion, and (ii) the pixel to pixel differences of the bias. The first correction was planned to be performed using the mean and standard deviation of all available pixels of the pre-scan columns, which contain no photo-electrons nor thermally excited electrons. However, in-flight margins turned out to be unusable, and this correction is now done using a constant whose value depends on the reading mode and is provided by the instrument team. The second correction is carried out using images taken with zero exposure time during ground-based calibration of the CCD or during the commissioning of the instrument.

Gain—Each exposure is corrected using the model derived from the on-ground characterisation of the CCD.

Linearisation—The number of ADU counts for a given pixel is not a linear function of the number of photo-electrons, in particular outside the linear zone. The deviation from linearity has been measured from a series of repeated exposures with different exposure times acquired during the ground-based characterization of the CCD and the calibration of the instrument. The DRP is ensuring the non-linearity correction even for on-board stacked images.

Dark current—The characterisation of the CCD in the laboratory has shown that the pixel-to-pixel dark current is very low, below the read-out noise for a one-minute exposure. In the pre-launch version of the DRP, the correction was dynamically evaluated using the dedicated blind pixels provided by the 32 dark columns: the median value was the dark current, and the dispersion the read-out noise. However, to cope with the in-flight instrument and to account for the difference in read-out time between the first and the last read-out, the dark current is simply corrected using a constant weighted by the effective integration time of the pixels.

Flat-field—To correct the variation in pixel-to-pixel response in the best manner, CHEOPS uses chromatic flat fields adapted to the spectral energy distribution of the target. These chromatic flat fields are built as the weighted linear combination of monochromatic flat fields at wavelengths from 400 nm to 1100 nm acquired in the laboratory (see §1.1.1.19). In the red part where fringing is observed in the flat fields, the broad-band / filter is used instead of the narrow-band flats. For a given target, the weights are calculated from the expected values of the target spectral energy distribution (SED) at each of these wavelengths and associated bandwidth, taking into account the throughput and quantum efficiency. The SED is derived through a collection of 40 stellar SED ranging from 3700 K to 7200 K. In addition, there is the possibility to use a simple black body distribution. The correction has the flexibility to accommodate for flat-field updates from the in-flight monitoring and characterisation activities.

Orbital events flagging—The low-altitude orbit induces some events that are fully predictable and directly impact the quality of the data and thus the duty cycle of the observations. This is the case of the South Atlantic Anomaly (SAA) crossing, which results in a high increase in the number of high energy protons impacts. The Earth illumination also has an influence on the satellite, especially during the transition from light to penumbra/shadow and the reverse producing temperature variations. Depending on the pointing direction, the Sun, Earth and Moon reflected light can affect the data, resulting in a significant degradation of the duty cycle. The minimal avoidance angles for an exposure to be valid are 120° for the Sun, 35° for the Earth, and 15° for the Moon. The events flagging module is thus in charge of identifying the images that were acquired in poor observing conditions. It is done using an orbital model reconstruction and following housekeeping temperatures.

Geometric model—This model is used to transform the coordinates of the target star into (x,y) coordinates in the CCD frames and vice versa.

1.1.1.24 Correction

Inputs to this module are the calibrated images produced by the previous steps in photo-electron counts. These individual calibrated frames are corrected in this module of the pipeline from environmental effects. It includes the following successive processing steps:

- the bad pixels flagging and correction;
- the cosmic ray hits flagging and correction;
- the smear estimate and correction;

- the background and stray light estimate and correction.

A description of the main science algorithms used in the correction module are presented below.

Smearing—The telescope has no shutter, and thus the CCD remains exposed while the transfer process is in progress. As a consequence, at each readout, vertical transfer smear trails are generated. In most configurations, the target itself can be the source of smearing trails, as seen in Figure 29. However, bright stars on the CCD, even outside the subarray, also produce their own smearing trails.




Figure 29: Example of smearing trails produced by the bright target and a nearby star outside the subarray.

In a standard smear correction procedure, for a frame k , the smear estimate in a given pixel is done using the average value of the top overscan pixels, divided by the number of pixels in a column of the whole CCD. This usual correction, however, assumes a static image which is not the case for CHEOPS. Indeed, the field rotation combined with the jitter makes the field on the CCD slightly different between two consecutive exposures. Consequently, the charge in a given pixel will sweep across two different images during the transfer process to the storage area (For more details, see Hoyer et al., 2020). In its in-flight version, the DRP corrects smearing trails from a series of simulated images that correspond to each image of the visit. These simulations allow which frames have to be corrected to be identified as well as how to correct them. They are built including all stars within the full CHEOPS field of view, down to magnitude 19 in the GAIA bandpass, and scaled to the CHEOPS flux using an updated PSF.

Bad pixels—The bad pixels module detects and corrects pixels that give a temporary or a permanent abnormal response. Three different kinds of bad pixels are considered by the pipeline, which identified them through different algorithms:

- Hot/dead pixels, i.e., pixels that are damaged and whose flux sensitivity is abnormal.
- Random telegraphic signals (RTS) due to discrete changes of the pixel behaviour caused, for instance, by irradiation or fluctuations in the current.
- Cosmic ray hits on the CCD with energy deposit over a strip range of pixels, depending on the angle of incidence.

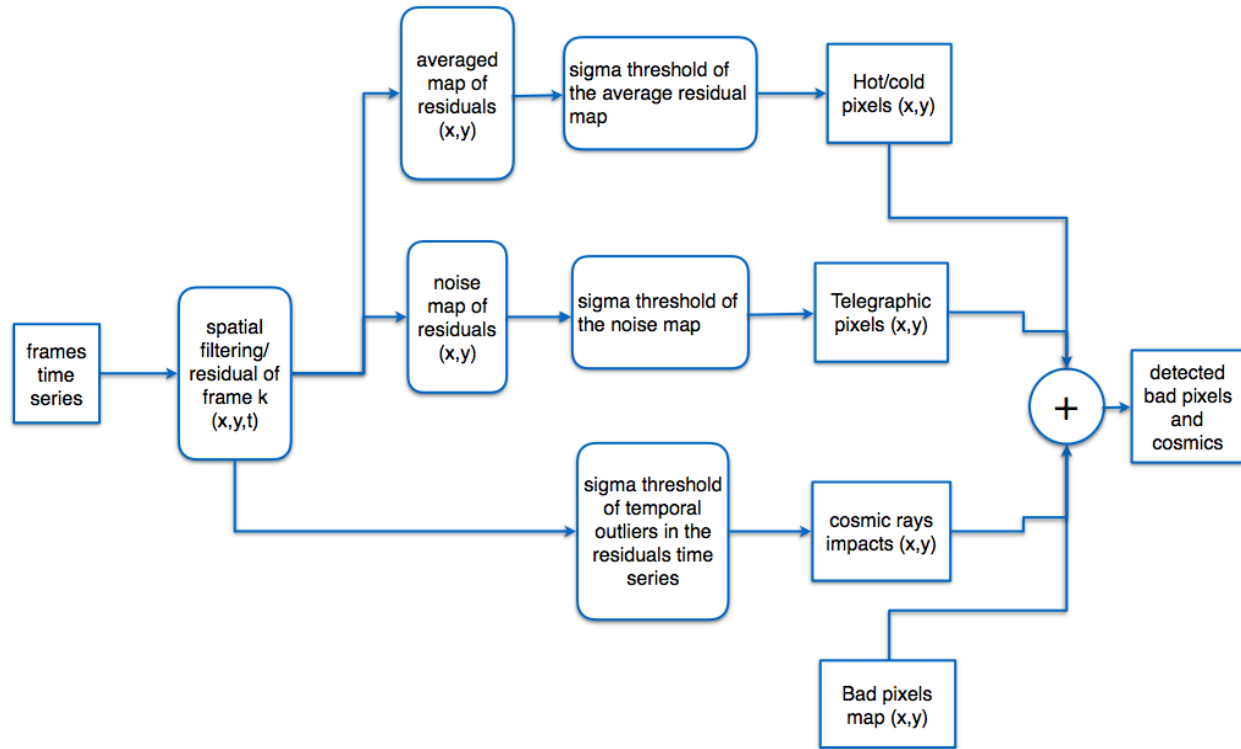


Figure 30: Flow chart for the bad pixel detection algorithm. Note the bad pixels map is one of the input reference files.

Whatever their nature, bad pixels are searched for and identified in the complete subarray, that is not only the pixels of the target PSF pixels but also those of the background. Figure 30 presents a flow chart describing the bad pixels identification algorithm. The outlier identification is done on the image time series. To remove the variations due to field motions and stellar signal, each image is filtered with a 2D median 10×10 pixels filter and the residuals are then calculated as the difference between the time series and the filtered one. Two maps are then created: one is the time average of each residual, the second one is its dispersion in time. On the first map, positive and negative outliers are identified as hot/cold pixels by applying a threshold to the absolute value. Random telegraphic signals are identified from the second map as noisy pixels. Cosmic ray hits, which are sudden events that affect a single image in the series, are identified through a time-varying estimate of the residuals for each pixel. At the very last step, this dynamic bad pixel events identification is merged with the map that has been built from the characterisation of the CCD in the laboratory. This map has to be updated with in-orbit characterisation and monitoring of the

instrument. As a result of the detection step, a bad pixel map of the sub-array is built, with all pixels flagged according to their status, which currently are:

- 2: Dead;
- 4: Hot;
- 8: Telegraphic.

Currently, Bad Pixels are not corrected by the DRP. Only detected cosmic rays are corrected by interpolation of adjacent pixels using a 2-D cubic spline.

Applying a threshold for outliers detection is not straightforward in the case of CHEOPS because (i) the noise is not stationary due to the permanent roll, and (ii) shot noise is dominated by the jitter noise at pixel level, especially near the centre of the target PSF, which exhibit sharp gradients. To mitigate possible issues with long integration exposures, short exposure *imagettes* are downloaded to increase the cosmic ray contrast. These imagettes are 60 pixels in diameter, centred on the target, made of sub-stacks of one or more readouts. They are processed by the pipeline the same way as the regular frames but without the dark, flat fielding, and background correction as they are not supposed to be used for photometry purposes but for cosmic ray identification by DRP. For each imagette, a cosmic ray map is built. The sub-stack imagettes are then put in the same reference frame so that to identify which pixels have been flagged, as imagettes could be affected by the jitter. Once a pixel in the target PSF is identified as being impacted by cosmic rays (through inspection of the relevant imagettes), its value on the sub-array is replaced by its cosmic ray-cleaned pixel value, and the associated bad pixels map is updated accordingly. Finally, the statistics of bad pixel detection per frame/image is given in detail in the DRP report.

Background and straylight correction— These two components are not evaluated separately but as a single quantity labelled as “background”. Estimation of the total background component does not rely on the interpolation of the flux of dark windows, as usually implemented for high precision photometry, because of the field rotation due to the telescope roll over its orbit. Indeed, taking into account the rotation would involve a continuous adjustment of the windows position, generating additional noise. Instead, a mask is generated on the sub-array which excludes the corners, in case of square image, and the central target (Figure 31, left). It follows the target centroid and is large enough so that to ensure that (i) it remains insensitive to the field motion, and (ii) it maximises the number of pixels used for the background flux estimate. The latter is done by calculating the flux value corresponding to the maximum of the flux distribution of all pixels inside the mask (**Error! Reference source not found.**Figure 31)The corresponding background light curve is built and provided to the user.

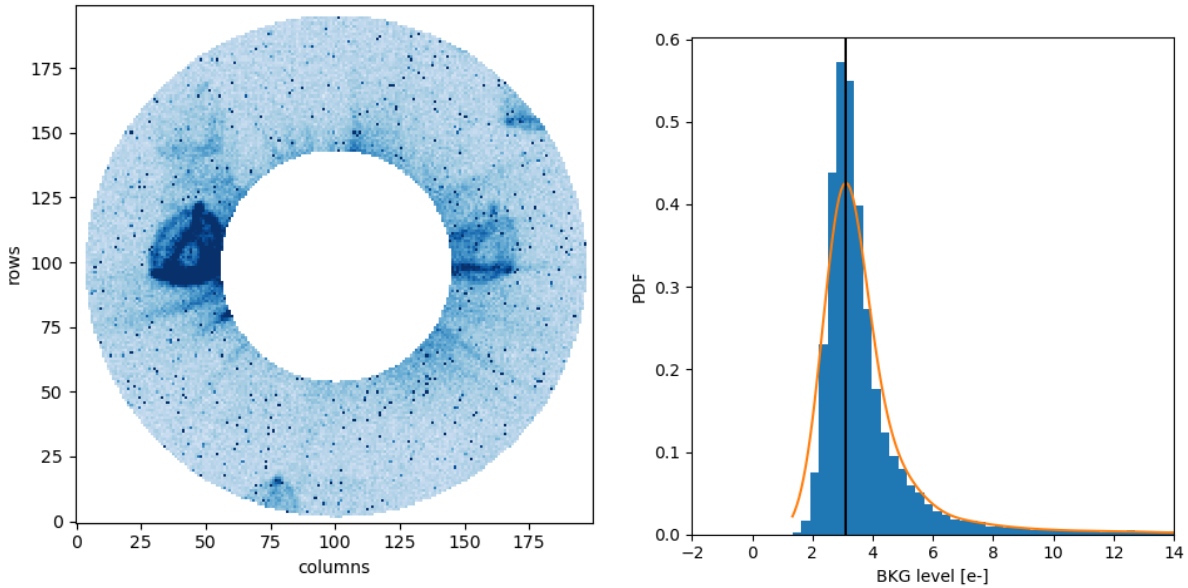


Figure 31: (Left) Example of a mask used for the background estimation. The target and the corners of the square sub-array are excluded from the mask. (Right) Histogram of the flux of all the pixels encompassed by the background mask (blue) and the best skewed Gaussian fit. Possible contaminating stars will contribute to the tail of the distribution.

1.1.1.25 Photometry

The last step in the data flow is the extraction of the photometric signal of the target. It is done using an aperture photometry method. The latter uses a binary mask centred on the target photocenter to follow the pointing jitter. The border of the mask is properly weighted at subpixel level. For all targets, the pipeline provides the light curve derived using a default aperture of 25 pixels which takes into account the shape of the in-flight PSF. However, as the size of the aperture that optimises the signal-to-noise ratio depends in fact on a number of factors such as the target magnitude and the local crowding, three additional light curves, derived with different aperture radii, are also provided.

To define the optimal aperture, simulated CHEOPS images of the field-of-view corresponding to the sub-array are generated (Figure 32). These images include not only the target but also all stars that are in the CHEOPS star catalogue provided to the DRP and encompassed in the sub-array. They are built considering the roll angle, to follow the instantaneous motion of the platform. The PSF is calculated using the reference PSF, the flux being scaled according to the CHEOPS magnitude of the stars and the exposure time of the frames, using Vega as reference for the magnitude system. From these images, the pipeline computes the signal-to-noise ratio as a function of the aperture radius,

$$Signal_to_Noise = F_{target} / \sqrt{F_{target} + F_{cont} + \sigma_{cont}^2 + N_{readout} \times N_{pix} \times RON^2}$$

where

- F_{target} is the median of the simulated flux of the target within the aperture;
- F_{cont} and σ_{cont} are the median and standard deviation of the simulated flux of background stars;
- the last term corresponds to the read-out noise within the aperture. It is estimated as $N_{pix} N_{stack} RON^2$, where
 - o N_{pix} is the number of pixels encompassed by the aperture;
 - o $N_{readout}$ is the number of individual read-out frames in the science image;
 - o RON is the read-out noise in units of electrons per pixel.

The radius at which the SNR is maximum is then automatically determined (Figure 33). It corresponds to the so-called *optimal* aperture. This optimal aperture is not focused on the RMS criteria of the light curve but accounts for the effect of close contaminating stars, which might affect the transit depth in the case of crowded fields. Two other light curves are also provided to give the user an idea of how the RMS of the light curve evolves with the aperture size.




Figure 32: An image of a star observed by CHEOPS (left) and the of simulated field-of-view calculated in the pipeline using all stars in the catalogue (middle). The right image shows the simulated image with the contaminants only . Note that for the latter image it is set on a different scale for clarity.

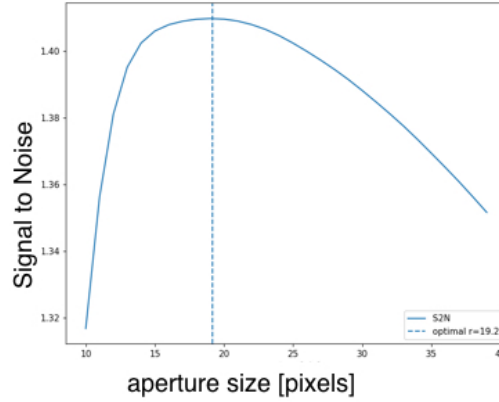


Figure 33: The SNR as a function of the aperture size. The vertical dashed line indicates the optimal aperture size that maximises the SNR.

Because of the non-symmetric shape of the PSF and of the rotation of the platform, a contaminant close to the target could have part of its flux entering and exiting the aperture. As a consequence, the *contamination factor* is not necessarily constant but can vary in time (Figure 34). A contamination light curve is thus also derived from the simulated images for each of the light curves that are provided to the users.




Figure 34: (Top) Target and contaminants flux normalized to the average of the target flux. There is a clear modulation with the rotation period of the platform. (Bottom) Flux of the contaminants only. The modulation is due to contaminants only.

Data reduction products, report and reference files

From the processing of a given visit, the DRP issues several products that the user will retrieve from the archive. In particular, there are four light curves calculated with different apertures size:

- the 25-pixel, default aperture;
- the optimal aperture based on the optimisation of the SNR;
- the 22.5-pixel, that is the lower limit of the aperture not affected by jitter noise;
- the 30-pixel, which includes 120% of the PSF flux.

Each of these light curves comes with its own contamination curve. In addition to these final products, the user will have the possibility to get additional by-products of the processing, such as, for example, bad pixel maps, the background light curve, and smearing light curves.



Data Reduction Report

May 04, 2018

Contents

1	Final Light curve	1
1.1	Observation summary	2
1.2	Metrics summary	3
1.3	Noise curve	3
1.4	Pipeline status	3
2	Summary of the processing stages	4
3	Process breakdown	5
3.1	Bias correction	5
3.2	ADU to electron conversion	6
3.3	Dark current correction	7
3.4	Flat field correction	7
3.5	Smear correction	8
3.6	Bad pixels correction	9
3.7	Background correction	10
3.8	Centroids estimation	11
3.9	Flux correlations	12
3.10	Contamination Estimation	12
3.11	Aperture Optimization	13

Figure 35: First page of one Data Reduction Report that will be automatically generated for the user each time a visit will be processed.

Finally, the DRP generates a report for the user (Figure 35). This report provides a summary of the observations carried out in a given visit but also of the CHEOPS DRP used for the data reduction. It also shows the evolution of the data through the pipeline from the raw images to the final reduced light curves. For each step of the processing, some figures and the derived light curves are shown with a measure of the dispersion of the target light curve before and after. This will allow the user to follow what treatment has been done on the data, the quality of the

processing and of the final result. Examples of data reduction report can be found with the simulated data sets presented in Section 0.

Finally, and for reference, the DRP uses the following reference files:

- EXT_APP_DE1
- EXT_PRE_StarCatalogue
- REF_APP_BadPixelMap
- REF_APP_BiasFrame
- REF_APP_DarkFrame
- REF_APP_FlatField
- REF_APP_WhitePSF
- REF_APP_OversampledWhitePSF
- REF_APP_CCDLinearisation
- REF_APP_GainCorrection
- REF_APP_Throughput
- REF_APP_SED
- REF_APP_PixelScale
- REF_APP_QE
- REF_APP_VisitConstraints
- REF_APP_Limits-ScienceHk
- REF_APP_ObtReset

A detailed description of these reference files with all header keywords is provided in the Data Product Definition Document which is available from the call webpage,

What the DRP does not correct for (aka, systematic errors)

During the first months of CHEOPS observations, the DRP evolved to accommodate real in-flight data. The last major module currently turned off is the bad pixel correction. So far it has been estimated to introduce more error than is introduced by the hot pixels. It is possible that it will be activated in the next DRP major release. Please consult the release notes contained in each DRP report.


Finally, the instrument will be monitored and, if relevant, aging corrections will be implemented.

Evolution of the DRP

The DRP is in constant evolution to further improve CHEOPS data quality and science output. It will soon receive another major release, DRP v14.0, upon which all existing data will be reprocessed, more apertures will become available, and some keywords will change.

To stay up to date with all DRP developments as well as other in-orbit updates for observers, please regularly consult:

→ <https://www.cosmos.esa.int/web/cheops-quest-observers-programme/in-orbit-updates>



End of document. Thank you for reading.