EUROPEAN SOUTHERN OBSERVATORY

Organisation Européenne pour des Recherches Astronomiques dans l'Hémisphère Austral

Europäische Organisation für astronomische Forschung in der südlichen Hemisphäre

VERY LARGE TELESCOPE

Reflex NIRPS Tutorial

Issue 3.3.13

Date 2025-07-02

Prepared:	A. Modigliani, V. D. Ivanov	2025-07-02	
	Name	Date	Signature
Approved:	W.Freudling		
	Name	Date	Signature
Released:	M. Sterzik		
	Name	Date	Signature

This page was intentionally left blank

Change record

Issue/Rev.	Date	Section/Parag. affected	Reason/Initiation/Documents/Remarks
3.3.13	02/07/2025	none	Public release

This page was intentionally left blank

Contents

1 Introduction to `Esoreflex`

This document is a tutorial designed to enable the user to to reduce his/her data with the ESO pipeline run under an user-friendly environmet, called `EsoReflex`, concentrating on high-level issues such as data reduction quality and signal-to-noise (S/N) optimisation.

`EsoReflex` is the ESO Recipe Flexible Execution Workbench, an environment to run ESO VLT pipelines which employs a workflow engine to provide a real-time visual representation of a data reduction cascade, called a workflow, which can be easily understood by most astronomers. The basic philosophy and concepts of Reflex have been discussed by Freudling et al. (2013A&A...559A..96F). Please reference this article if you use Reflex in a scientific publication.

Reflex and the data reduction workflows have been developed by ESO and instrument consortia and they are fully supported. If you have any issue, please have a look to https://support.eso.org to see if this has been reported before or open a ticket for further support.

A workflow accepts science and calibration data, as downloaded from the archive using the CalSelector tool[1] (with associated raw calibrations) and organises them into DataSets, where each DataSet contains one science object observation (possibly consisting of several science files) and all associated raw and static calibrations required for a successful data reduction. The data organisation process is fully automatic, which is a major time-saving feature provided by the software. The DataSets selected by the user for reduction are fed to the workflow which executes the relevant pipeline recipes (or stages) in the correct order. Full control of the various recipe parameters is available within the workflow, and the workflow deals automatically with optional recipe inputs via built-in conditional branches. Additionally, the workflow stores the reduced final data products in a logically organised directory structure employing user-configurable file names.

The NIRPS instrument is described in Wildi et al. (2017).

[1] *http://www.eso.org/sci/archive/calselectorInfo.html*

2 The NIRPS Workflow

The NIRPS workflow canvas is organised into a number of areas. From top-left to top-right you will find general workflow instructions, directory parameters, and global parameters. In the middle row you will find five boxes describing the workflow general processing steps in order from left to right, and below this the workflow actors themselves are organised following the workflow general steps.

2.1 Workflow Canvas Parameters

The workflow canvas displays a number of parameters that may be set by the user. Under "Setup Directories" the user is only required to set the `RAW_DATA_DIR` to the working directory for the dataset(s) to be reduced, which, by default, is set to the directory containing the demo data. The `RAW_DATA_DIR` is recursively scanned by the `Data Organiser` actor for input raw data. The directory `CALIB_DATA_DIR`, which is by default within the pipeline installation directory, is also scanned by the `Data Organiser` actor to find any static calibrations that may be missing in your dataset(s). If required, the user may edit the directories `BOOKKEEPING_DIR`, `LOGS_DIR`, `TMP_PRODUCTS_DIR`, and `END_PRODUCTS_DIR`, which correspond to the directories where book-keeping files, logs, temporary products and end products are stored, respectively (see the Reflex User Manual for further details; Forchì (2012)).

There is a mode of the `Data Organiser` that skips the built-in data organisation and uses instead the data organisation provided by the CalSelector tool. To use this mode, click on `Use CalSelector associations` in the `Data Organiser` properties and make sure that the input data directory contains the XML file downloaded with the CalSelector archive request (note that this does not work for all instrument workflows).

Under the "Global Parameters" area of the workflow canvas, the user may set the `FITS_VIEWER` parameter to the command used for running his/her favourite application for inspecting FITS files. Currently this is set by default to `fv`, but other applications, such as `ds9`, `skycat` and `gaia` for example, may be useful for inspecting image data. Note that it is recommended to specify the full path to the visualization application (an alias will not work).

By default the `EraseDirs` parameter is set to `false`, which means that no directories are cleaned before executing the workflow, and the recipe actors will work in Lazy Mode (see Section 5.4.2), reusing the previous pipeline recipe outputs if input files and parameters are the same as for the previous execution, which saves considerable processing time. Sometimes it is desirable to set the `EraseDirs` parameter to `true`, which forces the workflow to recursively delete the contents of the directories specified by `BOOKKEEPING_DIR`, `LOGS_DIR`, and `TMP_PRODUCTS_DIR`. This is useful for keeping disk space usage to a minimum and will force the workflow to fully re-reduce the data each time the workflow is run.

The parameter `RecipeFailureMode` controls the behaviour in case that a recipe fails. If set to `Continue`, the workflow will trigger the next recipes as usual, but without the output of the failing recipe, which in most of the cases will lead to further failures of other recipes without the user actually being aware of it. This mode might be useful for unattended processing of large number of datasets. If set to `Ask`, a pop-up window will ask whether the workflow should stop or continue. This is the default. Alternatively, the `Stop` mode will stop the workflow execution immediately.

The parameter `GlobalPlotInteractivity` controls whether the interactive windows will appear for those windows which are *enabled* by default. The possible values are `true, false`. Take into account that some

windows are disabled in the default configuration and therefore are not affected by this parameter.

The parameter `ProductExplorerMode` controls whether the `ProductExplorer` actor will show its window or not. The possible values are `Enabled`, `Triggered`, and `Disabled`. `Enabled` opens the ProductExplorer GUI at the end of the reduction of each individual dataset. `Triggered` (default and recommended) opens the ProductExplorer GUI when all the selected datasets have been reduced. `Disabled` does not display the ProductExplorer GUI.

3 Software Installation

`Esoreflex` and the workflows can be installed in different ways: via package repositories, via the `install_esoreflex` script or manually installing the software tar files.

The recommended way is to use the package repositories if your operating system is supported. The pipelines and Reflex can be installed from the ESO `macports` repositories that support macOS platforms, the and the `rpm/yum` repositories that support Fedora and CentOS platforms. For any other operating system it is recommended to use the `install_esoreflex` script.

The installation from package repository requires administrative privileges (typically granted via sudo), as it installs files in system-wide directories under the control of the package manager. If you want a local installation, or you do not have sudo privileges, or if you want to manage different installations on different directories, then use the `install_esoreflex` script. Note that the script installation requires that your system fulfill several software prerequisites, which might also need sudo privileges.

Reflex 2.11.x needs java JDK 11 to be installed.

Please note that in case of major or minor (affecting the first two digit numbers) Reflex upgrades, the user should erase the `$HOME/KeplerData`, `$HOME/.kepler` directories if present, to prevent possible aborts (i.e. a hard crash) of the `esoreflex` process.

3.1 Installing `Esoreflex` workflows via `macports`

This method is supported for the macOS operating system. It is assumed that macports (*http://www.macports.org*) is installed. Please read the full documentation at *http://www.eso.org/sci/software/pipelines/installation/macports.html*, which also describes the versions of macOS that are currently supported.

3.2 Installing `Esoreflex` workflows via `rpm/yum/dnf`

This method is supported for Fedora and CentOS platforms and requires sudo rights. Please read the full documentation at *http://www.eso.org/sci/software/pipelines/installation/rpm.html*, which also describes the versions of Fedora and CentOS that are currently supported.

3.3 Installing `Esoreflex` workflows via `install_esoreflex`

This method is recommended for operating systems other than what indicated above, or if the user has no sudo rights. Software dependencies are not fulfilled by the installation script, therefore the user has to install all the prerequisites before running the installation script.

The software pre-requisites for `Reflex 2.11.5` may be found at:
http://www.eso.org/sci/software/pipelines/reflex_workflows

To install the `Reflex 2.11.5` software and demo data, please follow these instructions:

1. From any directory, download the installation script:

```
wget https://eso.org/sci/software/pipelines/install_esoreflex
```

2. Make the installation script executable:

```
chmod u+x install_esoreflex
```

3. Execute the installation script:

```
./install_esoreflex
```

   and the script will ask you to specify three directories: the download directory <download_dir>, the software installation directory <install_dir>, and the directory to be used to store the demo data <data_dir>. If you do not specify these directories, then the installation script will create them in the current directory with default names.

4. Follow all the script instructions; you will be asked whether to use your Internet connection (recommended: yes), the pipelines and demo-datasets to install (note that the installation will remove all previously installed pipelines that are found in the same installation directory).

5. To start Reflex, issue the command:

```
<install_dir>/bin/esoreflex
```

   It may also be desirable to set up an alias command for starting the Reflex software, using the shell command alias. Alternatively, the PATH variable can be updated to contain the <install_dir>/bin directory.

3.4 Demo Data

Together with the pipeline the user will also receive a demo data set, that allows you to run the Reflex NIRPS workflow without any changes in parameters. This way you have a data set to perform a run-time verification after the installation and to experiment with before you start to work on your own data.

Note that you will need a minimum of ∼21 GB, ∼2.9 GB and ∼23 GB of free disk space for the directories <download_dir>, <install_dir> and <data_dir>, respectively.

4 Quick Start: Reducing The Demo Data

For the user who is keen on starting reductions without being distracted by detailed documentation, we describe the steps to be performed to reduce the science data provided in the NIRPS demo data set supplied with the `esoreflex 2.11.5` release. By following these steps, the user should have enough information to perform a reduction of his/her own data without any further reading:

1. First, type:

   ```
   esoreflex -l
   ```

   If the `esoreflex` executable is not in your path, then you have to provide the command with the executable full path `<install_dir>/bin/esoreflex -l` . For convenience, we will drop the reference to `<install_dir>`. A list with the available `esoreflex` workflows will appear, showing the workflow names and their full path.

2. Open the `nirps` by typing:

   ```
   esoreflex nirps&
   ```

   Alternatively, you can type only the command `esoreflex` the empty canvas will appear (Figure 4.0.1) and you can select the workflow to open by clicking on `File -> Open File`. Note that the loaded workflow will appear in a new window. The `nirps` workflow is shown in Figure 4.0.2.

3. To aid in the visual tracking of the reduction cascade, it is advisable to use component (or actor) highlighting. Click on `Tools -> Animate at Runtime`, enter the number of milliseconds representing the animation interval (100 ms is recommended), and click `OK`.

4. Change directories set-up. Under "Setup Directories" in the workflow canvas there are seven parameters that specify important directories (green dots).

   By default, the `ROOT_DATA_DIR`, which specifies the working directory within which the other directories are organised. is set to your `$HOME/reflex_data` directory. All the temporary and final products of the reduction will be organized under sub-directories of `ROOT_DATA_DIR`, therefore make sure this parameter points to a location where there is enough disk space. To change `ROOT_DATA_DIR`, double click on it and a pop-up window will appear allowing you to modify the directory string, which you may either edit directly, or use the `Browse` button to select the directory from a file browser. When you have finished, click `OK` to save your changes.

   Changing the value of `RAW_DATA_DIR` is the only necessary modification if you want to process data other than the demo data

5. Click the  button to start the workflow

6. The workflow will highlight the `Data Organiser` actor which recursively scans the raw data directory (specified by the parameter `RAW_DATA_DIR` under "Setup Directories" in the workflow canvas) and constructs the datasets. Note that the raw and static calibration data must be present either

in `RAW_DATA_DIR` or in `CALIB_DATA_DIR`, otherwise datasets may be incomplete and cannot be processed. However, if the same reference file was downloaded twice to different places this creates a problem as `esoreflex` cannot decide which one to use.

7. The `Data Set Chooser` actor will be highlighted next and will display a "Select Datasets" window (see Figure 4.0.3) that lists the datasets along with the values of a selection of useful header keywords[2]. The first column consists of a set of tick boxes which allow the user to select the datasets to be processed. By default all complete datasets which have not yet been reduced will be selected. A full description of the options offered by the `Data Set Chooser` will be presented in Section 6.3.2.

8. Click the `Continue` button and watch the progress of the workflow by following the red highlighting of the actors. A window will show which dataset is currently being processed.

9. Once the reduction of all datasets has finished, a pop-up window called *Product Explorer* will appear, showing the datasets which have been reduced together with the list of final products. This actor allows the user to inspect the final data products, as well as to search and inspect the input data used to create any of the products of the workflow. Figure 4.0.4 shows the Product Explorer window. A full description of the *Product Explorer* will be presented in Section 6.3.5.

10. After the workflow has finished, all the products from all the datasets can be found in a directory under `END_PRODUCTS_DIR` named after the workflow start timestamp. Further subdirectories will be found with the name of each dataset.

Well done! You have successfully completed the quick start section and you should be able to use this knowledge to reduce your own data. However, there are many interesting features of `Reflex` and the NIRPS workflow that merit a look at the rest of this tutorial.

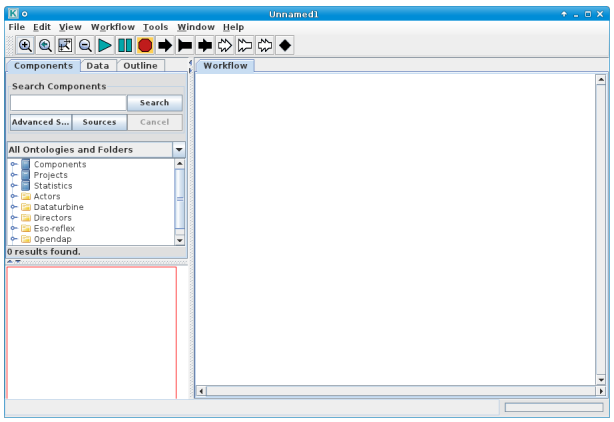

Figure 4.0.1: *The empty* `Reflex` *canvas.*

[2]The keywords listed can be changed by double clicking on the `DataOrganiser` Actor and editing the list of keywords in the second line of the pop-up window. Alternatively, instead of double-clicking, you can press the right mouse button on the `DataOrganiser` Actor and select `Configure Actor` to visualize the pop-up window.

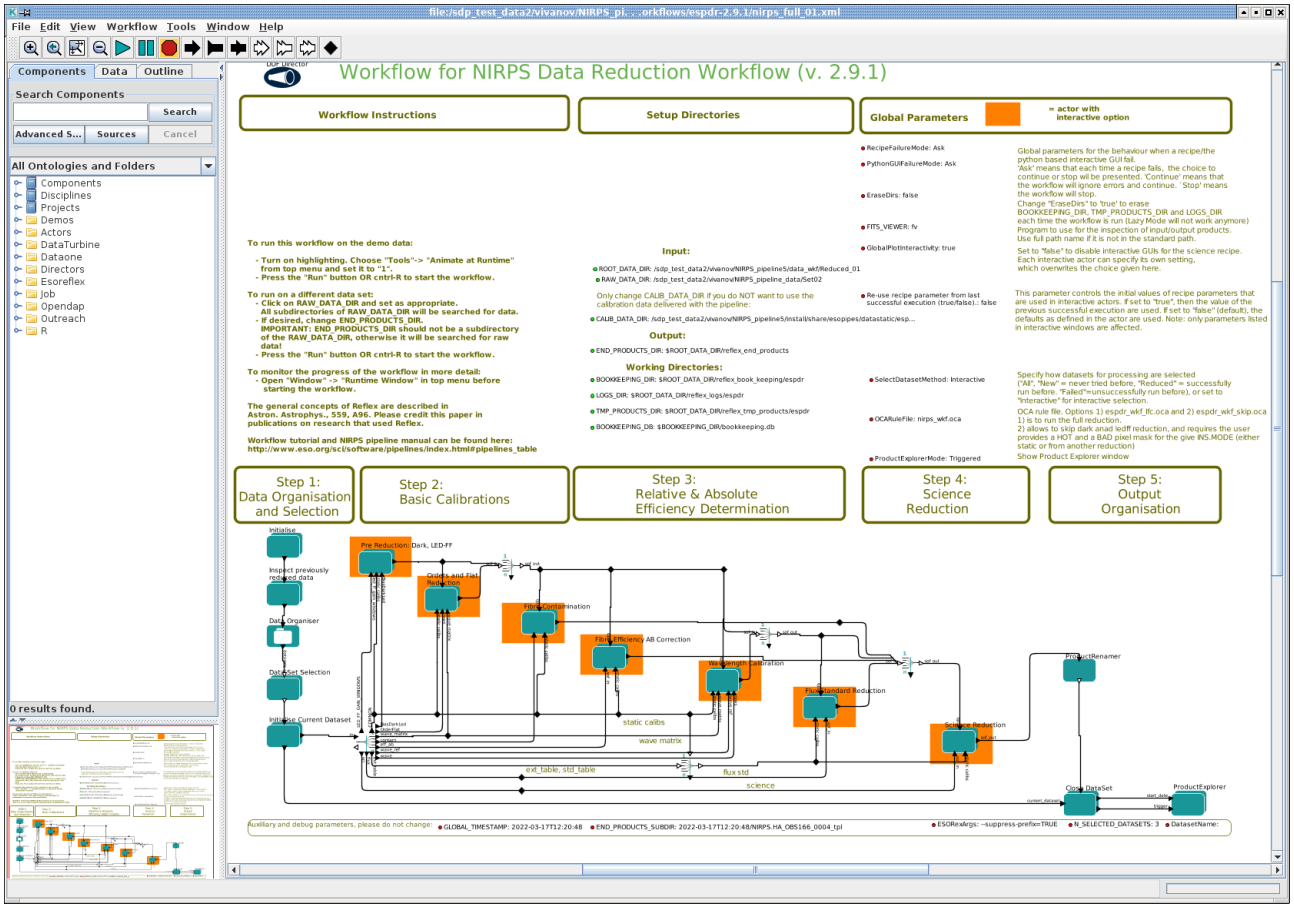

Figure 4.0.2: *NIRPS workflow general layout. Newer version may differ slightly.*

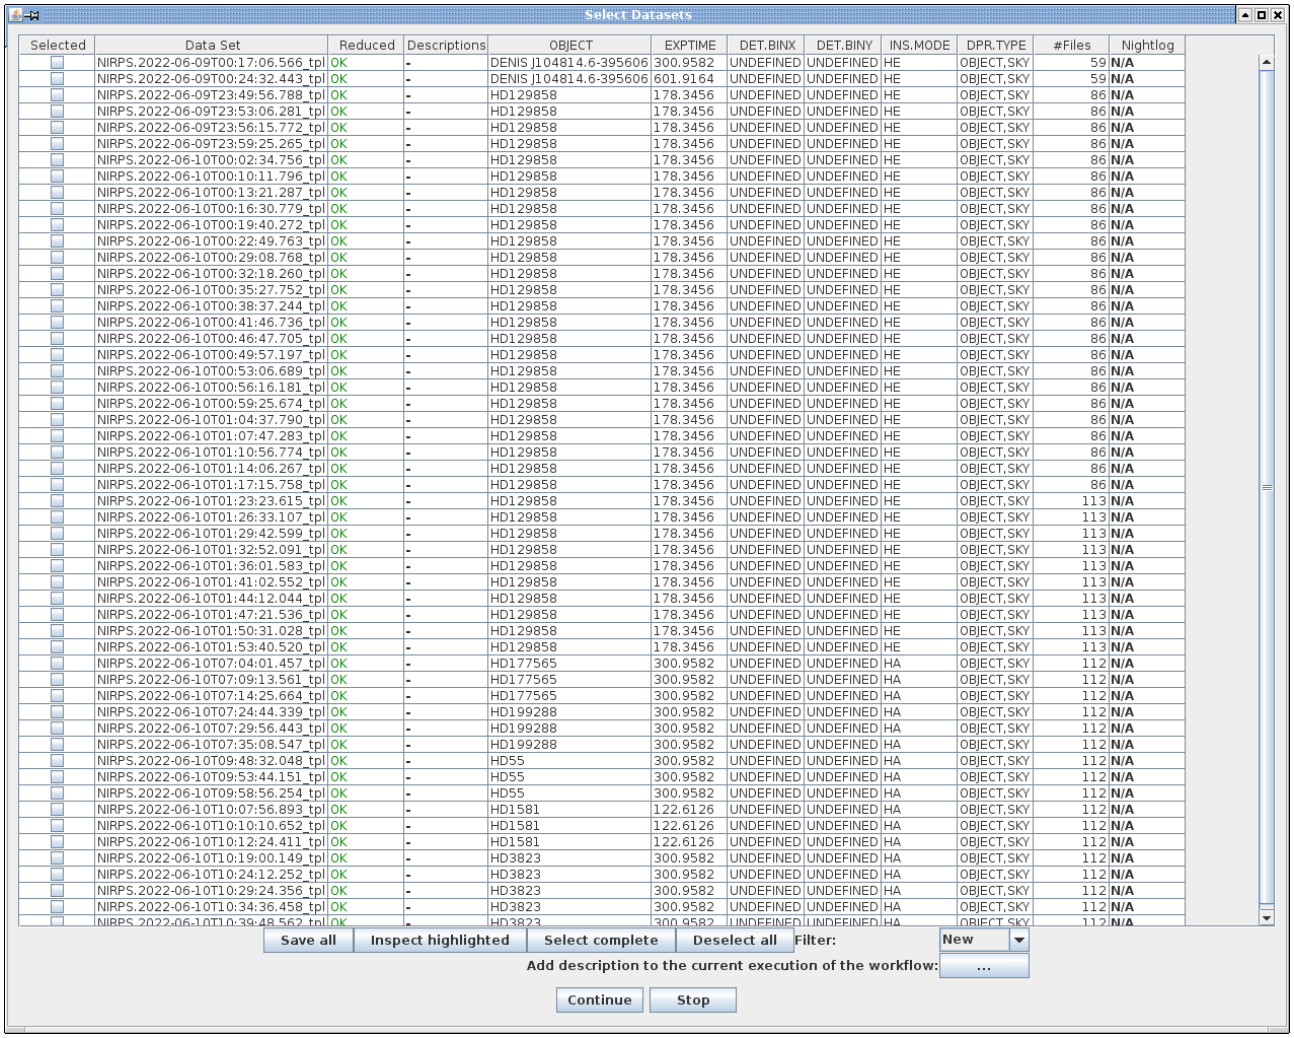

Figure 4.0.3: *The "Select Datasets" pop-up window.*

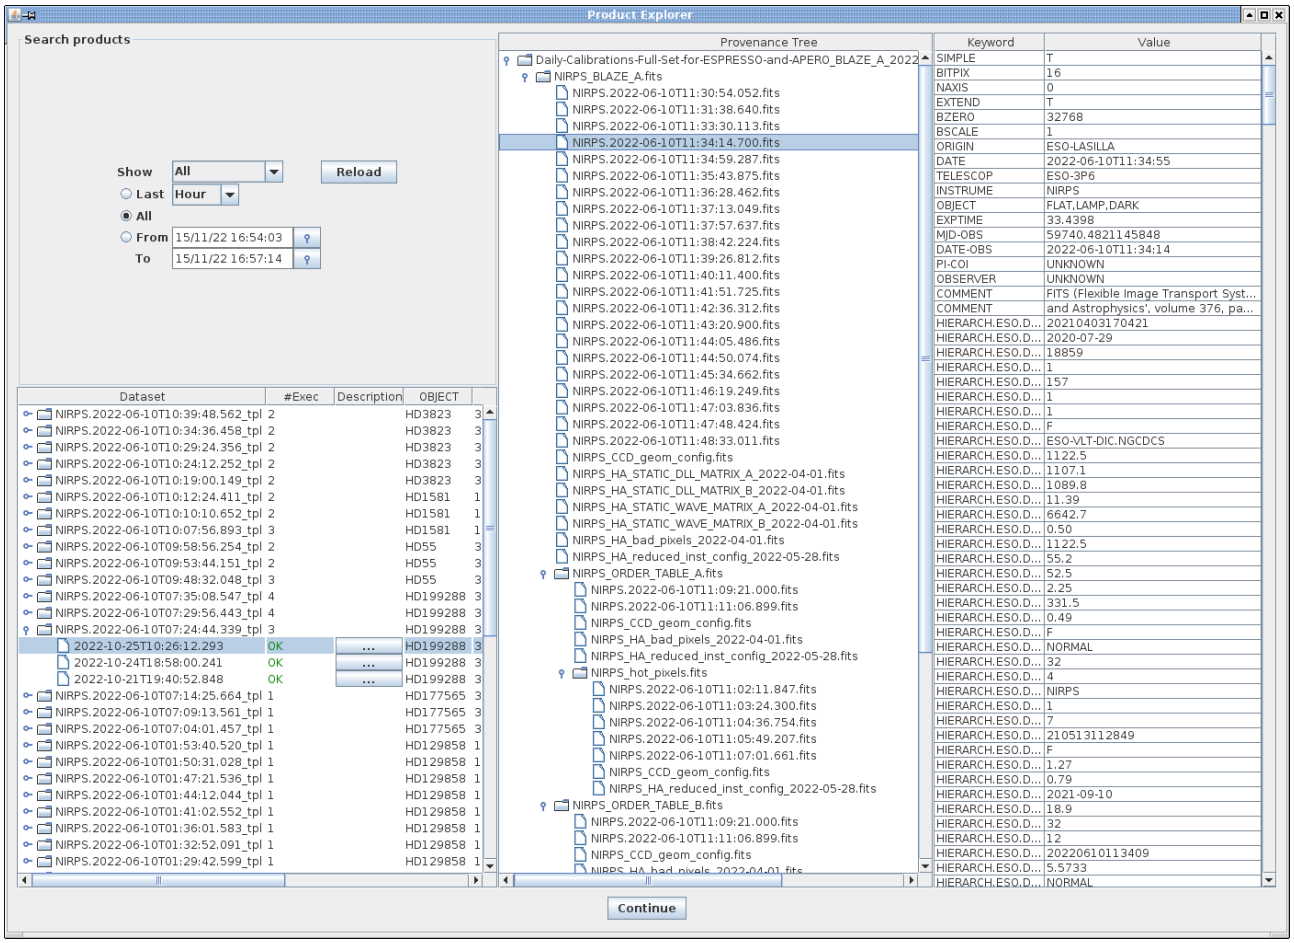

Figure 4.0.4: *The Provenance Explorer shows all datasets reduced in previous executions together with the full reduction chain for all the pipeline products.*

5 About the main `esoreflex` canvas

5.1 Saving And Loading Workflows

In the course of your data reductions, it is likely that you will customise the workflow for various data sets, even if this simply consists of editing the `ROOT_DATA_DIR` to a different value for each data set. Whenever you modify a workflow in any way, you have the option of saving the modified version to an `XML` file using `File -> Export As` (which will also open a new workflow canvas corresponding to the saved file). The saved workflow may be opened in subsequent `esoreflex` sessions using `File -> Open`. Saving the workflow in the default Kepler format (.kar) is only advised if you do not plan to use the workflow with another computer.

5.2 Buttons

At the top of the `esoreflex` canvas are a set of buttons which have the following functions:

-  - Zoom in.
-  - Reset the zoom to 100%.
-  - Zoom the workflow to fit the current window size (Recommended).
-  - Zoom out.
-  - Run (or resume) the workflow.
-  - Pause the workflow execution.
-  - Stop the workflow execution.

The remainder of the buttons (not shown here) are not relevant to the workflow execution.

5.3 Workflow States

A workflow may only be in one of three states: executing, paused, or stopped. These states are indicated by the yellow highlighting of the  ,  , and  buttons, respectively. A workflow is executed by clicking the  button. Subsequently the workflow and any running pipeline recipe may be stopped immediately by clicking the  button, or the workflow may be paused by clicking the  button which will allow the current actor/recipe to finish execution before the workflow is actually paused. After pausing, the workflow may be resumed by clicking the  button again.

5.4 Workflow Actors

5.4.1 Simple Actors

Simple actors have workflow symbols that consist of a single (rather than multiple) green-blue rectangle. They may also have an icon within the rectangle to aid in their identification. The following actors are simple actors:

- 

- The `DataOrganiser` actor.
- 

- The `DataSetChooser` actor (inside a composite actor).
- 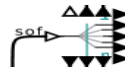

- The `FitsRouter` actor Redirects files according to their categories.
- 

- The `ProductRenamer` actor.
- 

- The `ProductExplorer` actor (inside a composite actor).

Access to the parameters for a simple actor is achieved by right-clicking on the actor and selecting `Configure Actor`. This will open an "Edit parameters" window. Note that the `Product Renamer` actor is a jython script (Java implementation of the Python interpreter) meant to be customised by the user (by double-clicking on it).

5.4.2 Lazy Mode

By default, all `RecipeExecuter` actors in a pipeline workflow are "Lazy Mode" enabled. This means that when the workflow attempts to execute such an actor, the actor will check whether the relevant pipeline recipe has already been executed with the same input files and with the same recipe parameters. If this is the case, then the actor will not execute the pipeline recipe, and instead it will simply broadcast the previously generated products to the output port. The purpose of the Lazy Mode is therefore to minimise any reprocessing of data by avoiding data re-reduction where it is not necessary.

One should note that the actor's Lazy Mode depends on the contents of the directory specified by the parameter `BOOKKEEPING_DIR` and the relevant FITS file checksums. Any modification to the directory contents and/or the file checksums will cause the corresponding actor to run the pipeline recipe again when executed, thereby re-reducing the input data.

The re-reduction of data at each execution may sometimes be desirable. To force a re-reduction of data for any single `RecipeExecuter` actor in the workflow, right-click the actor, select `Configure Actor`, and

uncheck the Lazy mode parameter tick-box in the "Edit parameters" window that is displayed. For many workflows the `RecipeExecuter` actors are actually found inside the composite actors in the top level workflow. To access such embedded `RecipeExecuter` actors you will first need to open the sub-workflow by right-clicking on the composite actor and then selecting `Open Actor`.

To force the re-reduction of all data in a workflow (i.e. to disable Lazy mode for the whole workflow), you must uncheck the Lazy mode for every single `RecipeExecuter` actor in the entire workflow. It is also possible to change the name of the bookkeeping directory, instead of modifying any of the Lazy mode parameters. This will also force a re-reduction of the given dataset(s). A new reduction will start (with the lazy mode still enabled), but the results of previous reduction will not be reused. Alternatively, if there is no need to keep any of the previously reduced data, one can simply set the `EraseDirs` parameter under the "Global Parameters" area of the workflow canvas to `true`. This will then remove all previous results that are stored in the bookkeeping, temporary, and log directories before processing the input data, in effect, starting a new clean data reduction and re-processing every input dataset. *Note: The option* `EraseDirs = true` *does not work in* `esoreflex` *version 2.9.x and makes the workflow to crash.*

6 Reducing your own data

In this section we describe how to reduce your own data set.

First, we suggest the reader to familiarize with the workflow by reducing the demo dataset first (Section 4), but it is not a requirement.

6.1 The `esoreflex` command

We list here some options associated to the `esoreflex` command. We recommend to try them to familiarize with the system. In the following, we assume the `esoreflex` executable is in your path; if not you have to provide the full path `<install_dir>/bin/esoreflex`

To see the available options of the `esoreflex` command type:

```
esoreflex -h
```

The output is the following.

```
-h | -help              print this help message and exit.
-v | -version           show installed Reflex version and pipelines and exit.
-l | -list-workflows    list available installed workflows and from
                        ~/KeplerData/workflows.
-n | -non-interactive   enable non-interactive features.
-e | -explore           run only the Product Explorer in this workflow
-p <workflow> | -list-parameters <workflow>
                        lists the available parameters for the given
                        workflow.
-config <file>          allows to specify a custom esoreflex.rc configuration
                        file.
-create-config <file>   if <file> is TRUE then a new configuration file is
                        created in ~/.esoreflex/esoreflex.rc. Alternatively
                        a configuration file name can be given to write to.
                        Any existing file is backed up to a file with a '.bak'
                        extension, or '.bakN' where N is an integer.
-debug                  prints the environment and actual Reflex launch
                        command used.
```

6.2 Launching the workflow

We list here the recommended way to reduce your own datasets. Steps 1 and 2 are optional and one can start from step 3.

1. Type: `esoreflex -n <parameters> nirps` to launch the workflow non interactively and reduce all the datasets with default parameters.

<parameters> allows you to specify the workflow parameters, such as the location of your raw data and the final destination of the products.

For example, type (in a single command line):

```
esoreflex -n
  -RAW_DATA_DIR /home/user/my_raw_data
  -ROOT_DATA_DIR /home/user/my_reduction
  -END_PRODUCTS_DIR $ROOT_DATA_DIR/reflex_end_products
  nirps
```

to reduce the complete datasets that are present in the directory /home/user/my_raw_data and that were not reduced before. Final products will be saved in /home/user/my_reduction/reflex_end_products, while book keeping, temporary products, and logs will be saved in sub-directories of /home/user/my_reduction/. If the reduction of a dataset fails, the reduction continues to the next dataset. It can take some time, depending on the number of datasets present in the input directory. For a full list of workflow parameters type esoreflex -p nirps. Note that this command lists only the parameters, but does not launch the workflow.

Once the reduction is completed, one can proceed with optimizing the results with the next steps.

2. Type:

   esoreflex -e nirps

   to launch the Product Explorer. The Product Explorer allows you to inspect the data products already reduced by the nirps esoreflex workflow. Only products associated with the workflow default bookkeeping database are shown. To visualize products associated to given bookkeeping database, pass the full path via the BOOKKEEPING_DB parameter:

   esoreflex -e BOOKKEEPING_DB <database_path> nirps

   to point the product explorer to a given <database_path>, e.g., /home/username/reflex/reflex_bookkeeping/test.db

   The Product Explorer allows you to inspect the products while the reduction is running. Press the button [Refresh] to update the content of the Product Explorer. This step can be launched in parallel to step 1.

   A full description of the Product Explorer will be given in Section 6.3.5

3. Type:

   esoreflex nirps &

   to launch the nirps esoreflex workflow. The nirps workflow window will appear (Fig. 4.0.2). Please configure the set-up directories ROOT_DATA_DIR, RAW_DATA_DIR, and other workflow parameters as needed. Just double-click on them, edit the content, and press [OK]. Remember to specify the same <database_path> as for the Product Explorer, if it has been opened at step #2, to synchronize the two processes.

4. (Recommended, but not mandatory) On the main esoreflex menu set Tools –> Animate at Runtime to 1 in order to highlight in red active actors during execution.

5. Press the button  to start the workflow. First, the workflow will highlight and execute the Initialise actor, which among other things will clear any previous reductions if required by the user (see Section 2.1).

Secondly, if set, the workflow will open the Product Explorer, allowing the user to inspect previously reduced datasets (see Section 6.3.5 for how to configure this option).

6.3 Workflow Steps

6.3.1 Data Organisation And Selection

The `DataOrganiser` (DO) is the first crucial component of a Reflex workflow. The DO takes as input `RAW_DATA_DIR` and `CALIB_DATA_DIR` and it detects, classifies, and organises the files in these directories and any subdirectories. The output of the DO is a list of "DataSets". A DataSet is a special Set of Files (SoF). A DataSet contains one or several science (or calibration) files that should be processed together, and all files needed to process these data. This includes any calibration files, and in turn files that are needed to process these calibrations. Note that different DataSets might overlap, i.e. some files might be included in more than one DataSet (e.g., common calibration files).

A DataSet lists three different pieces of information for each of its files, namely 1) the file name (including the path), 2) the file category, and 3) a string that is called the "purpose" of the file. The DO uses the OCA[3] rules to find the files to include in a DataSet, as well as their categories and purposes. The file category identifies different types of files, and it is derived by information in the header of the file itself. A category could for example be `RAW_CALIBRATION_1`, `RAW_CALIBRATION_2` or `RAW_SCIENCE`, depending on the instrument. The purpose of a file identifies the reason why a file is included in a DataSet. The syntax is `action_1/action_2/action_3/ ... /action_n`, where each `action_i` describes an intended processing step for this file (for example, creation of a `MASTER_CALIBRATION_1` or a `MASTER_CALIBRATION_2`). The actions are defined in the OCA rules and contain the recipe together with all file categories required to execute it (and predicted products in case of calibration data). For example, a workflow might include two actions `action_1` and `action_2`. The former creates `MASTER_ CALIBRATION_1` from `RAW_CALIBRATION_1`, and the later creates a `MASTER_CALIBRATION_2` from `RAW_CALIBRATION_2`. The `action_2` action needs `RAW_CALIBRATION_2` frames and the `MASTER_ CALIBRATION_1` as input. In this case, these `RAW_CALIBRATION_1` files will have the purpose `action_ 1/action_2`. The same DataSet might also include `RAW_CALIBRATION_1` with a different purpose; irrespective of their purpose the file category for all these biases will be `RAW_CALIBRATION_1`.

The Datasets created via the `DataOrganiser` will be displayed in the `DataSet Chooser`. Here the users have the possibility to inspect the various datasets and decide which one to reduce. By default, DataSets that have not been reduced before are highlighted for reduction. Click either `Continue` in order to continue with the workflow reduction, or `Stop` in order to stop the workflow. A full description of the `DataSet Chooser` is presented in Section 6.3.2.

Once the `Continue` is pressed, the workflow starts to reduce the first selected DataSet. Files are broadcasted according to their purpose to the relevant actors for processing.

[3]OCA stands for OrganisationClassificationAssociation and refers to rules, which allow to classify the raw data according to the contents of the header keywords, organise them in appropriate groups for processing, and associate the required calibration data for processing. They can be found in the directory `<install_dir>/share/esopipes/<pipeline-version>/reflex/`, carrying the extension `.oca`. The variable `<install_dir>` depends on the operative system and installation procedure. For installation through `rpm`: `<install_dir>=/usr`; for installation through macport `<install_dir>=/opt/local`; for installation through the installation script install_esoreflex it depends on the path specified during installation, e.g. `<install_dir>=<specified_path>/install`

The categories and purposes of raw files are set by the DO, whereas the categories and purpose of products generated by recipes are set by the `RecipeExecuter`. The file categories are used by the `FitsRouter` to send files to particular processing steps or branches of the workflow (see below). The purpose is used by the `SofSplitter` and `SofAccumulator` to generate input SoFs for the `RecipeExecuter`. The `SofSplitter` and `SofAccumulator` accept several SoFs as simultaneous input. The `SofAccumulator` creates a single output SoF from the inputs, whereas the `SofSplitter` creates a separate output SoF for each purpose.

6.3.2 DataSetChooser

The `DataSetChooser` displays the DataSets available in the "Select Data Sets" window, activating vertical and horizontal scroll bars if necessary (Fig. 4.0.3).

Some properties of the DataSets are displayed: the name, the number of files, a flag indicating if it has been successfully reduced (a green OK), if the reduction attempts have failed or were aborted (a red FAILED), or if it is a new dataset (a black "-"). The column "Descriptions" lists user-provided descriptions (see below), other columns indicate the instrument set-up and a link to the night log.

Sometimes you will want to reduce a subset of these DataSets rather than all DataSets, and for this you may individually select (or de-select) DataSets for processing using the tick boxes in the first column, and the buttons `Deselect All` and `Select Complete` at the bottom, or configure the "Filter" field at the bottom left. Available filter options are: "New" (datasets not previously reduced will be selected), "Reduced" (datasets previously reduced will be selected), "All" (all datasets will be selected), and "Failed" (dataset with a failed or aborted reduction will be selected).

You may also highlight a single DataSet in blue by clicking on the relevant line. If you subsequently click on `Inspect Highlighted`, then a "Select Frames" window will appear that lists the set of files that make up the highlighted DataSet including the full filename[4], the file category (derived from the FITS header), and a selection tick box in the right column. The tick boxes allow you to edit the set of files in the DataSet which is useful if it is known that a certain calibration frame is of poor quality (e.g: a poor raw flat-field frame). The list of files in the DataSet may also be saved to disk as an `ASCII` file by clicking on `Save As` and using the file browser that appears.

By clicking on the line corresponding to a particular file in the "Select Frames" window, the file will be highlighted in blue, and the file FITS header will be displayed in the text box on the right, allowing a quick inspection of useful header keywords. If you then click on `Inspect`, the workflow will open the file in the selected FITS viewer application defined by the workflow parameter `FITS_VIEWER`.

To exit from the "Select Frames" window, click `Continue`.

To add a description of the reduction, press the button `...` associated with the field "Add description to the current execution of the workflow" at the bottom right of the Select Dataset Window; a pop up window will appear. Enter the desired description (e.g. "My first reduction attempt") and then press `OK`. In this way, all the datasets reduced in this execution, will be flagged with the input description. Description flags can be visualized in the `SelectFrames` window and in the `ProductExplorer`, and they can be used to identify different reduction strategies.

[4] keep the mouse pointer on the file name to visualize the full path name.

To exit from the "Select DataSets" window, click either `Continue` in order to continue with the workflow reduction, or `Stop` in order to stop the workflow.

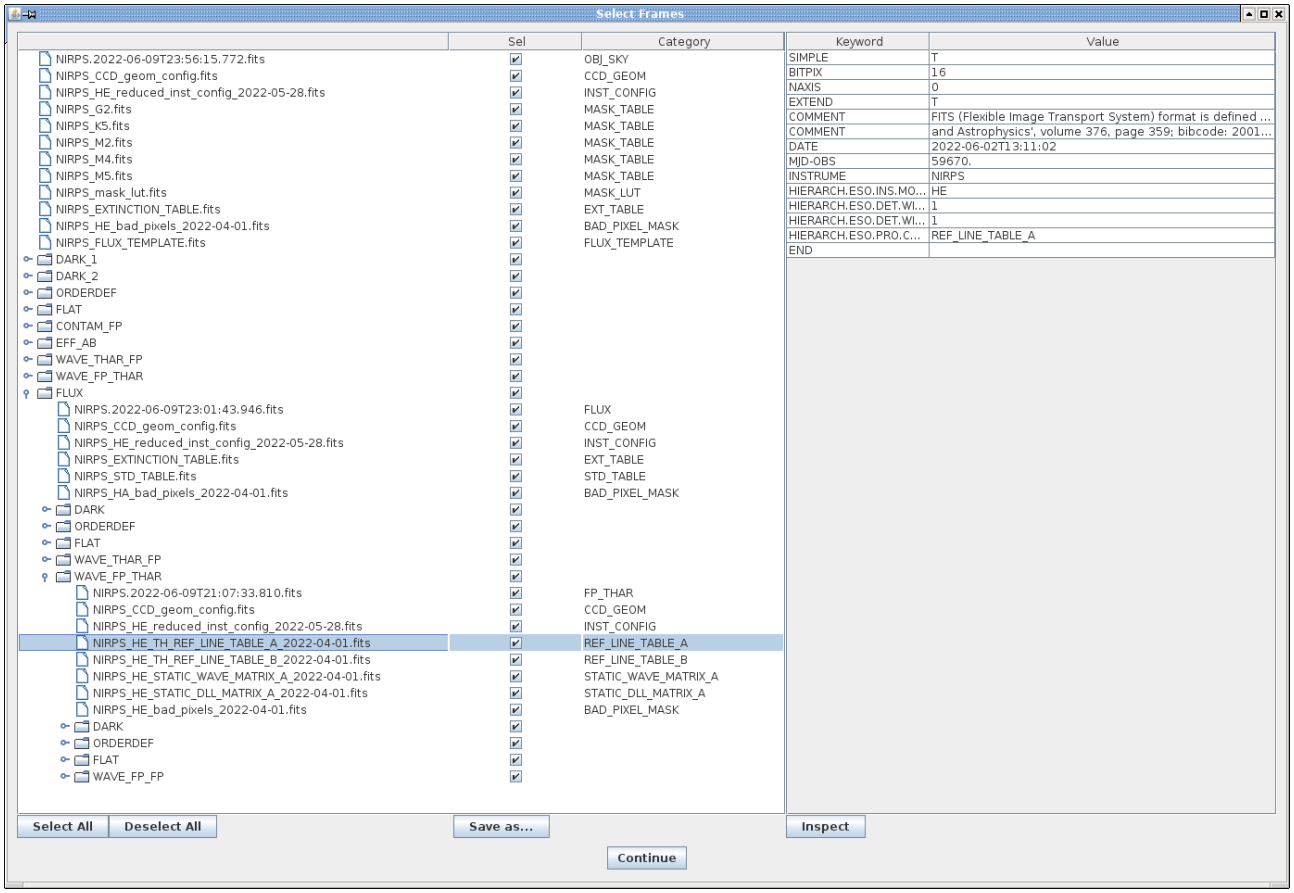

Figure 6.3.1: *The "Select Frames" window with a single file from the current Data Set highlighted in blue, and the corresponding FITS header displayed in the text box on the right. Hidden partially behind the "Select Frames" window is the "Select DataSets" window with the currently selected DataSet highlighted in blue.*

6.3.3 Recipe execution

Once the datasets to be reduced are selected, press the `Continue` button on the dataset organised window to proceed with the data reduction. The workflow will automatically execute the pipeline recipes and construct the `.sof` files to feed the pipeline recipes with. Each sof file will be saved in the `BOOKKEPING_DIR` directory (and subdirectory within it), depending on the recipe it is associated to and the execution time. The pipeline parameters can be changed as shown in figure 6.3.2.

6.3.4 Final products

Once a dataset is reduced (i.e. when the `espdr_sci_red` recipe is terminated), a window containing the list of science product pops up. Each file can be inspected with the selected fits viewer. Final science products will be stored in the `END_PRODUCTS_DIR`, and sorted by execution time and dataset identifier (i.e. the ARCHIVE file name of the science frame the dataset corresponds to). Default name for the science products is: `<TARG name>_S1D_FINAL_A_<DATE-OBS>.fits` Its category is _S1D_FINAL_A. It is a fits binary table, that contain the following columns:

WAVE	[Å]	Wavelength vector, in vacuum reference system.
FLUX	[erg cm^{-2} s^{-1} Å^{-1}] or [e–]	Flux of the final extracted spectrum in Fibre A, fully reduced till the end of the data reduction cascade. It is flux calibrated (if flux calibration was performed) and sky subtracted (if sky subtraction was performed).
ERR	[erg cm^{-2} s^{-1} Å^{-1}] or [e–]	Error associated to FLUX
QUAL		Quality flag associated to FLUX
SNR		Signal to noise ratio, obtained from FLUX/ERR.
WAVE_AIR	[Å]	Wavelength vector, in the air reference system.
FLUX_EL	[e–]	Extracted flux from fibre A, not sky subtracted nor flux calibrated.
ERR_EL	[e–]	Error associated to FLUX_EL.
QUAL_EL		Quality flag associated to FLUX_EL.
FLUX_CAL	[erg cm^{-2} s^{-1} Å^{-1}]	Extracted flux from fibre A, flux calibrated but not sky subtracted.
ERR_CAL	[erg cm^{-2} s^{-1} Å^{-1}]	Error associated to FLUX_CAL.
QUAL_CAL		Quality flag associated to FLUX_CAL.
FLUX_CAL_SKYSUB	[erg cm^{-2} s^{-1} Å^{-1}]	Extracted flux from fibre A, flux calibrated and sky subtracted.
ERR_CAL_SKYSUB	[erg cm^{-2} s^{-1} Å^{-1}]	Error associated to FLUX_CAL_SKYSUB.
QUAL_CAL_SKYSUB		Quality flag associated to FLUX_CAL_SKYSUB.
FLUX_EL_SKYSUB	[e–]	Extracted flux from fibre A, sky subtracted by not flux calibrated.
ERR_EL_SKYSUB	[e–]	Error associated to FLUX_EL_SKYSUB.
QUAL_EL_SKYSUB		Quality flag associated to FLUX_EL_SKYSUB.

Note: Some columns are present only if the corresponding step in the reduction cascade took place.

Other useful products are:
`<TARG name>_S1D_FINAL_B_<DATE-OBS>.fits`,
`<TARG name>_CCF_A_<DATE-OBS>.fits`,
`<TARG name>_CCF_B_<DATE-OBS>.fits`,
`<TARG name>_S1D_A_<DATE-OBS>.fits`,
`<TARG name>_S1D_B_<DATE-OBS>.fits`,
`<TARG name>_S1D_FLUXCAL_A_<DATE-OBS>.fits`,
`<TARG name>_S1D_SKYSUB_A_<DATE-OBS>.fits`,
`<TARG name>_S1D_SKYSUB_FLUXCAL_A_<DATE-OBS>.fits`,

<TARG name>_S2D_A_<DATE-OBS>.fits,
<TARG name>_S2D_B_<DATE-OBS>.fits,
<TARG name>_S2D_BLAZE_A_<DATE-OBS>.fits,
<TARG name>_S2D_BLAZE_A_<DATE-OBS>.fits,
<TARG name>_S2D_SKYSUB_A_<DATE-OBS>.fits,
<TARG name>_AIR_DLL_MATRIX_THAR_FP_A_<DATE-OBS>.fits,
<TARG name>_AIR_WAVE_MATRIX_THAR_FP_A_<DATE-OBS>.fits,
<TARG name>_BLAZE_A_<DATE-OBS>.fits,
<TARG name>_DLL_MATRIX_THAR_FP_A_<DATE-OBS>.fits and
<TARG name>_WAVE_MATRIX_THAR_FP_A_<DATE-OBS>.fits.

The user may find more details on the science recipe products in the NIRPS pipeline manual.

Note that we have added in the product filename the value of the FITS keyword <DATE-OBS> and that last five FITS products in the previous list, coming from calibration recipes, have been added to the END_PRODUCTS_DIR to ease the user is willing to analise the pipeline products with the NIRPS Data Analysis Software.

6.3.5 The ProductExplorer

The ProductExplorer is an interactive component in the esoreflex workflow whose main purpose is to list the final products with the associated reduction tree for each dataset and for each reduction attempt (see Fig. 4.0.4).

Configuring the ProductExplorer

You can configure the ProductExplorer GUI to appear after or before the data reduction. In the latter case you can inspect products as reduction goes on.

1. To display the ProductExplorer GUI at the end of the datareduction:

- Click on the global parameter “ProductExplorerMode” before starting the data reduction. A configuration window will appear allowing you to set the execution mode of the Product Explorer. Valid options are:
  - "Triggered" (default). This option opens the ProductExplorer GUI when all the selected datasets have been reduced.
  - "Enabled". This option opens the ProductExplorer GUI at the end of the reduction of each individual dataset.
  - “Disable”. This option does not display the ProductExplorer GUI.
- Press the  button to start the workflow.

2. To display the ProductExplorer GUI “before” starting the data reduction:

- double click on the composite Actor "Inspect previously reduced data". A configuration window will appear. Set to "Yes" the field "Inspect previously reduced data (Yes/No)". Modify the field "Continue reduction after having inspected the previously reduced data? (Continue/Stop/Ask)". "Continue" will continue the workflow and trigger the DataOrganizer. "Stop" will stop the workflow; "Ask" will prompt

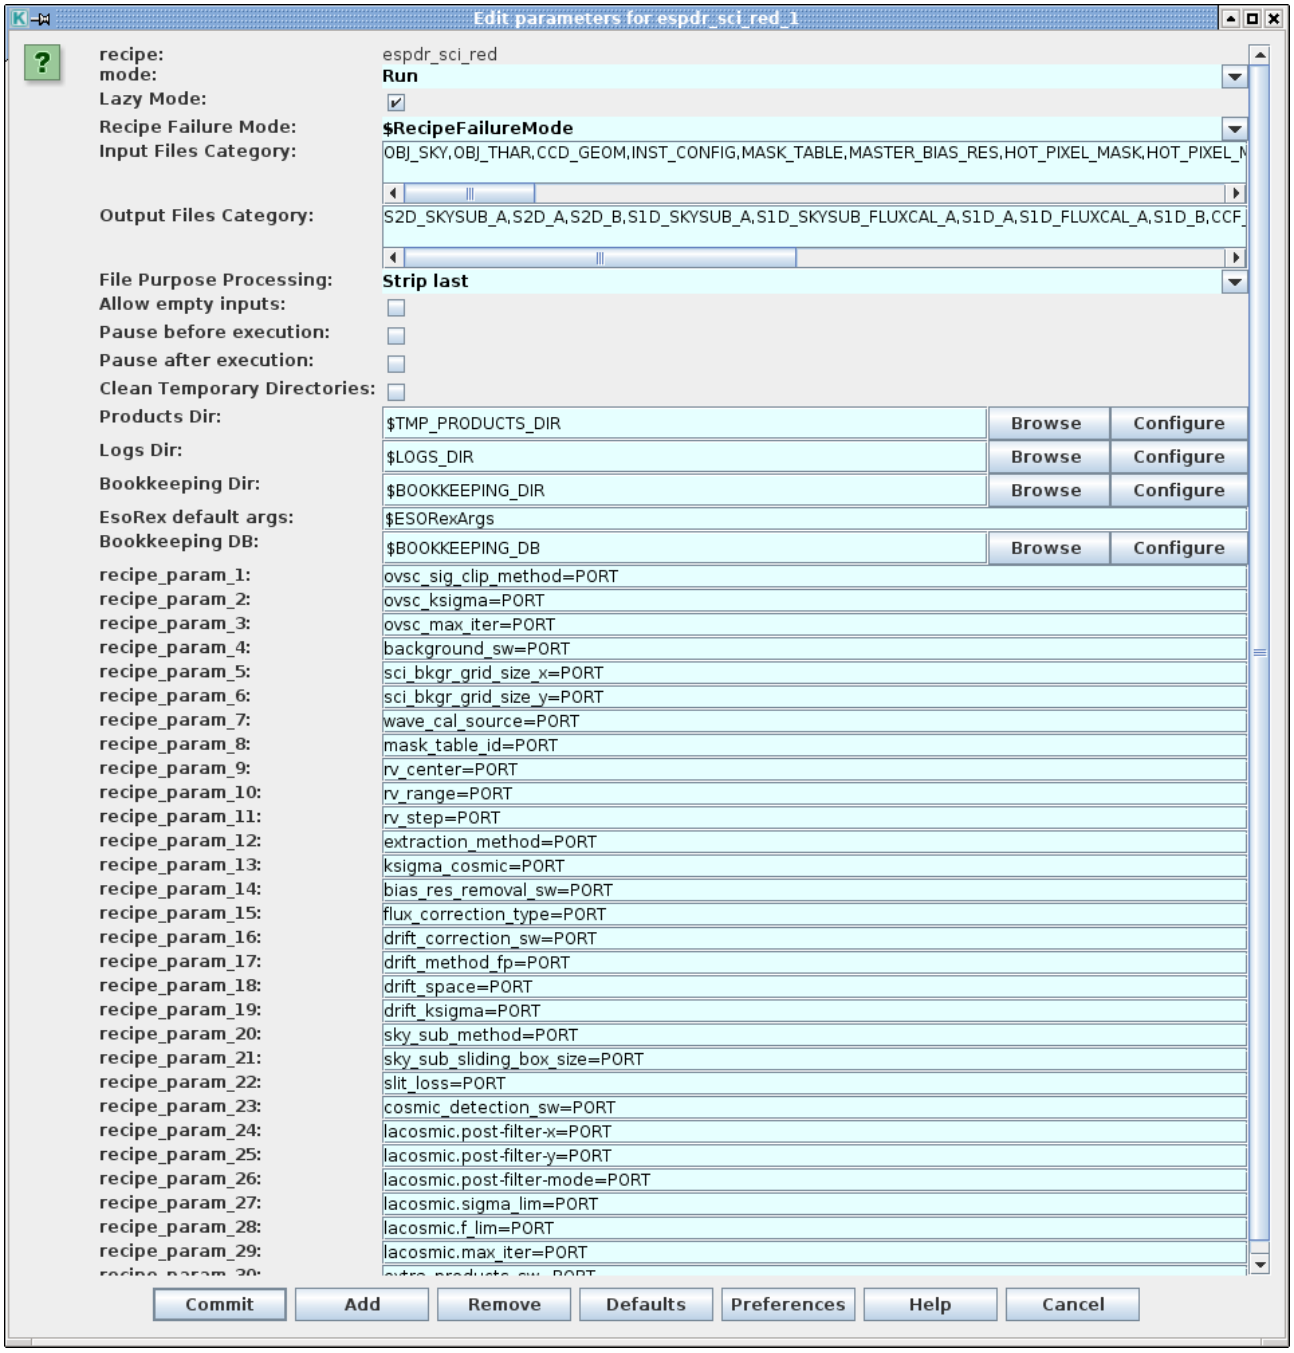

Figure 6.3.2: *The configuration parameters of a the science data reduction actor. The user may edit the configuration to change parameter values.*

another window deferring the decision whether continuing or not the reduction after having closed the Product Explorer.

- Press the  button to start the workflow. Now the ProductExplorer GUI will appear before starting the data organization and reduction.

Exploring the data reduction products

The left window of the ProductExplorer GUI shows the executions for all the datasets (see Fig. 4.0.4). Once you click on a dataset, you get the list of reduction attemps. Green and red flags identify successfull or unsucessfull reductions. Each reduction is linked to the “Description” tag assigned in the “Select Dataset” window.

1. To identify the desired reduction run via the “Description” tag, proceed as follows:

- Click on the symbol at the left of the dataset name. The full list of reduction attempts for that dataset will be listed. The column Exec indicates if the reduction was succesful (green flag: "OK") or not (red flag: "Failed").
- Click on the entries in the field "Description" to visualize the description you have entered associated to that dataset on the Select Dataset window when reducing the data.
- Identify the desired reduction run. All the products are listed in the central window, and they are organized following the data reduction cascade.

You can narrow down the range of datasets to search by configuring the field "Show" at the top-left side of the ProductExplorer (options are: "All", "Successful", "Unsuccessful"), and specifying the time range (Last, all, From-to).

2. To inspect the desired file, proceed as follows:

- Navigate through the data reduction cascade in the ProductExplorer by clicking on the files.
- Select the file to be inspected and click with the mouse right-hand button. The available options are:
  - Options available always:
    - Copy full path. It copies the full name of the file onto the clipboard. Shift+Ctr+v to past it into a terminal.
    - Inspect Generic. It opens the file with the fits viewer selected in the main workflow canvas.
    - Inspect with. It opens the file with an executable that can be specified (you have to provide the full path to the executable).
  - Options available for files in the `TMP_PRODUCTS_DIR` directory only:
    - command line. Copy of the environment configuration and recipe call used to generate that file.
    - Xterm. It opens an Xterm at the directory containing the file.
  - Options available for products associated to interactive windows only:
    - Display pipeline results. It opens the interactive windows associated to the recipe call that generated the file. Note that this is for visualization purposes only; the recipe parameters cannot be changed and the recipe cannot be re-run from this window.

7 Optimizing results through workflow Interactivity

This section describes the interactive GUIs provided in the NIRPS workflow. Some recipes provide interactive workflows. By default reflex interactivity is set to false for all calibration recipes and to true for the science data reduction. To activate calibration recipes interactivity the "GloalPlotInteractivityCalibs" reflex global parameter (you will see on the top left hand side of the reflex top level workflow layout) has to be set to true.

In general the reflex interactive GUI layout presents one image at the top, and one graph at the bottom. Specific functionality provided by buttons or sliders may appear in some interactive workflow as appropriate. The image at the top is used to display 2D images either in pixel-pixel or pixel-order (S2D format) scales. The graphs at the bottom are used to display scatter plots of quality control parameters. On the top left the user may change a few recipe parameters. To optimise image update time the images in pixel-pixel coordinates are resampled of a factor that depends on the detector binning, the larger the smaller the bin size. The following sections provide more details on each interactive workflow. Each section name refer to the main product or goal of the corresponding data reduction step.

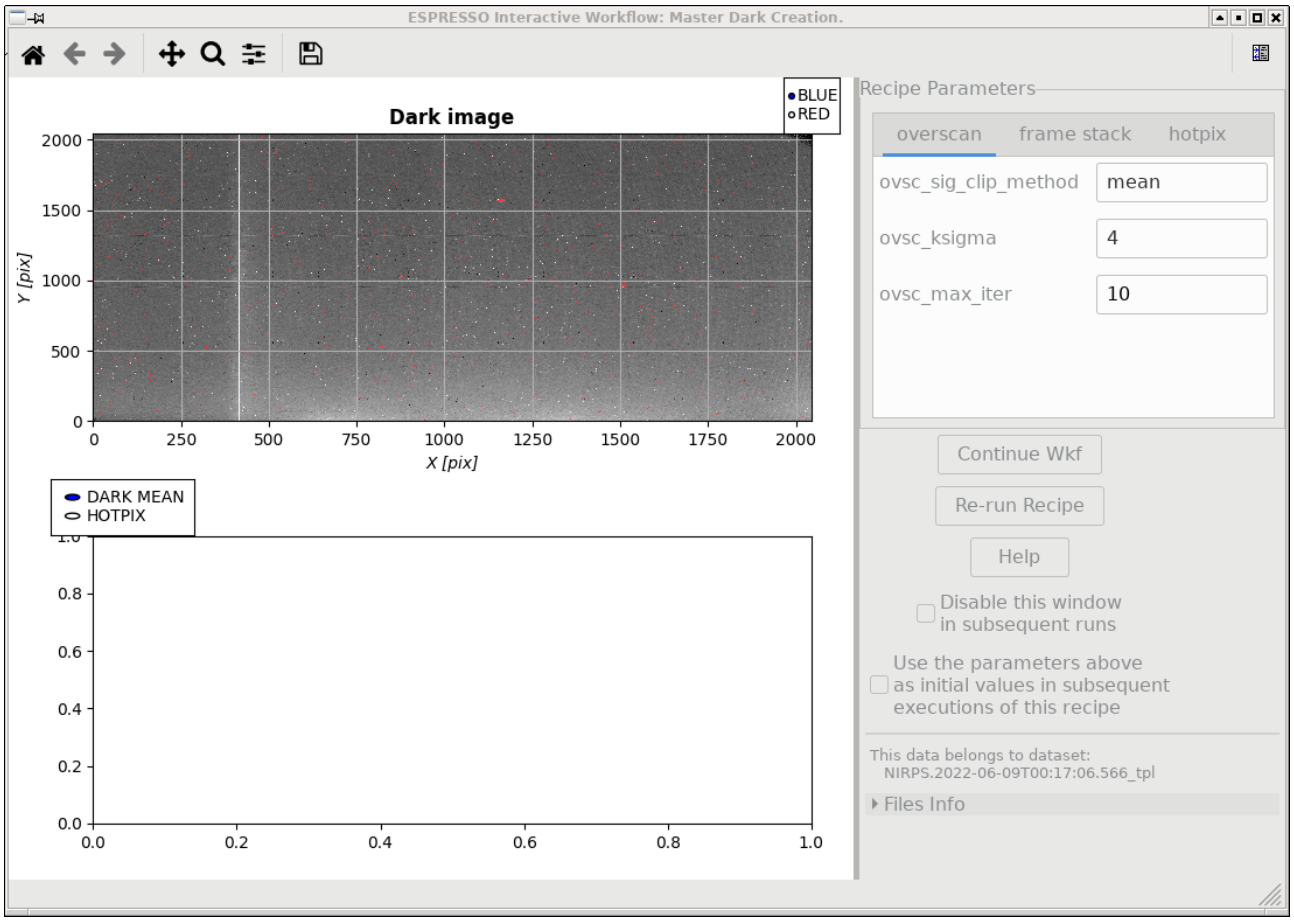

Figure 7.0.1: *Master Dark Interactive workflow.*

7.1 Master Dark

By default the master dark GUI displays at the top left an image of the master dark of the detector (see Fig. 7.0.1). Even though NIRPS has only one detector, for compatibility reasons the GUI offers options to display BLUE and RED detector, but always shows the same image. Pixels detected as hot are flagged with red in this panel (Fig. 7.1.1).

At the bottom left the GUI shows a scatter plot with the distribution of the mean of the DARK as a function of the read out port number. The user may select to change that and to plot the number of hot pixels as a function of the read out ports instead.

The GUI allows to change parameters controlling the frame stacking, the the hot pixels detection and the overscan. The latter is ignored for the infrared detector, but it was preserved for compatibility reasons.

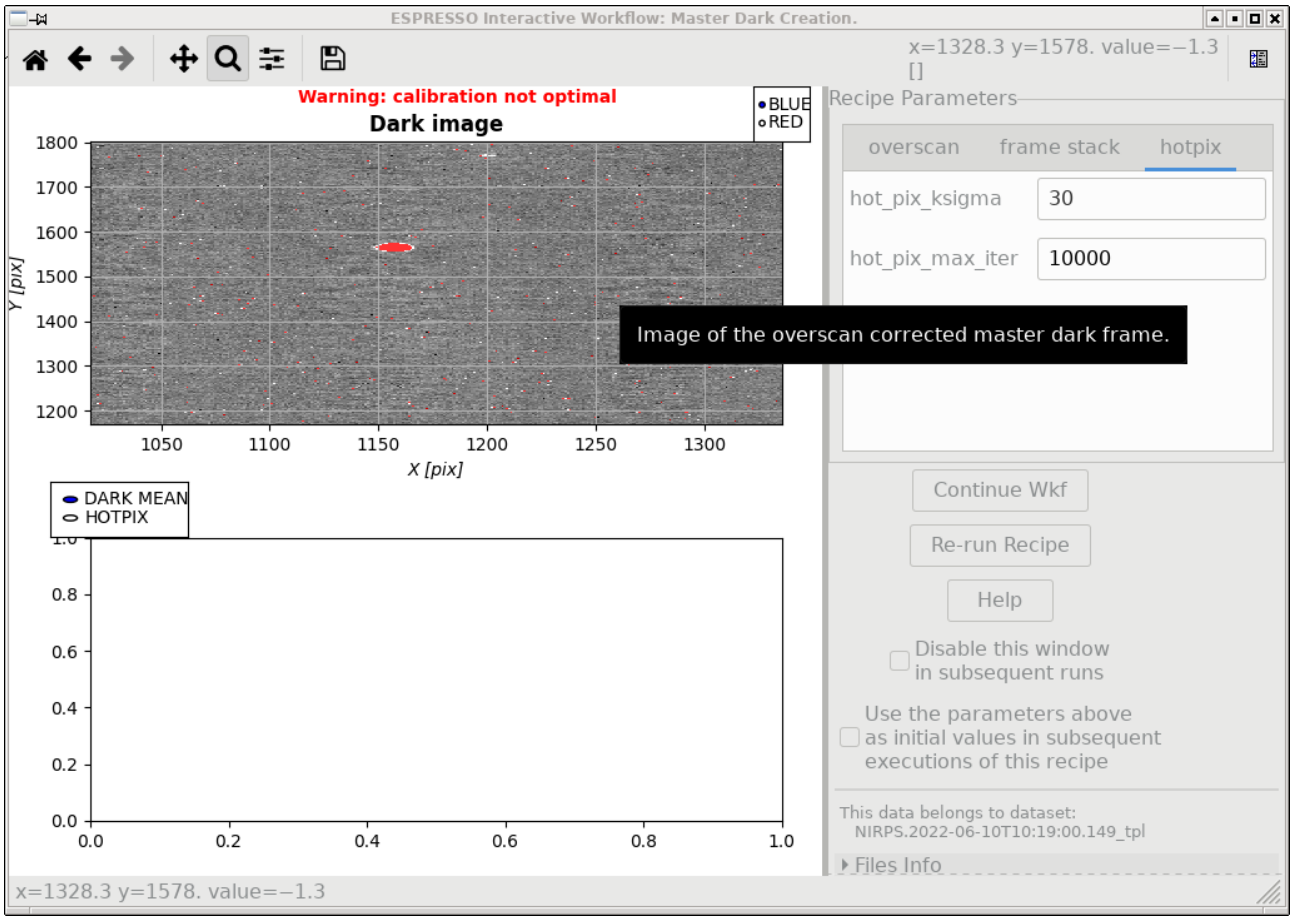

Figure 7.1.1: *Close up few centered on a cluster of hot pixels (marked with red).*

7.2 Orderdef

The orderdef GUI at the top left displays an image of the frame used to trace the orders with superimposing the computed order traces in green for a given detector (for NIRPS the choice between the BLUE and RED detectors is ignored) and a given fiber (Fig. 7.2.1).

At the bottom left a scatter plot is shown with a distribution of the standard deviation or the minimum or the maximum of the computed residuals as a function of the read-out port for the same choice of detector and fiber as for the top left panel.

As usual, the the GUI provides at the top right options to modify the recipe parameters.

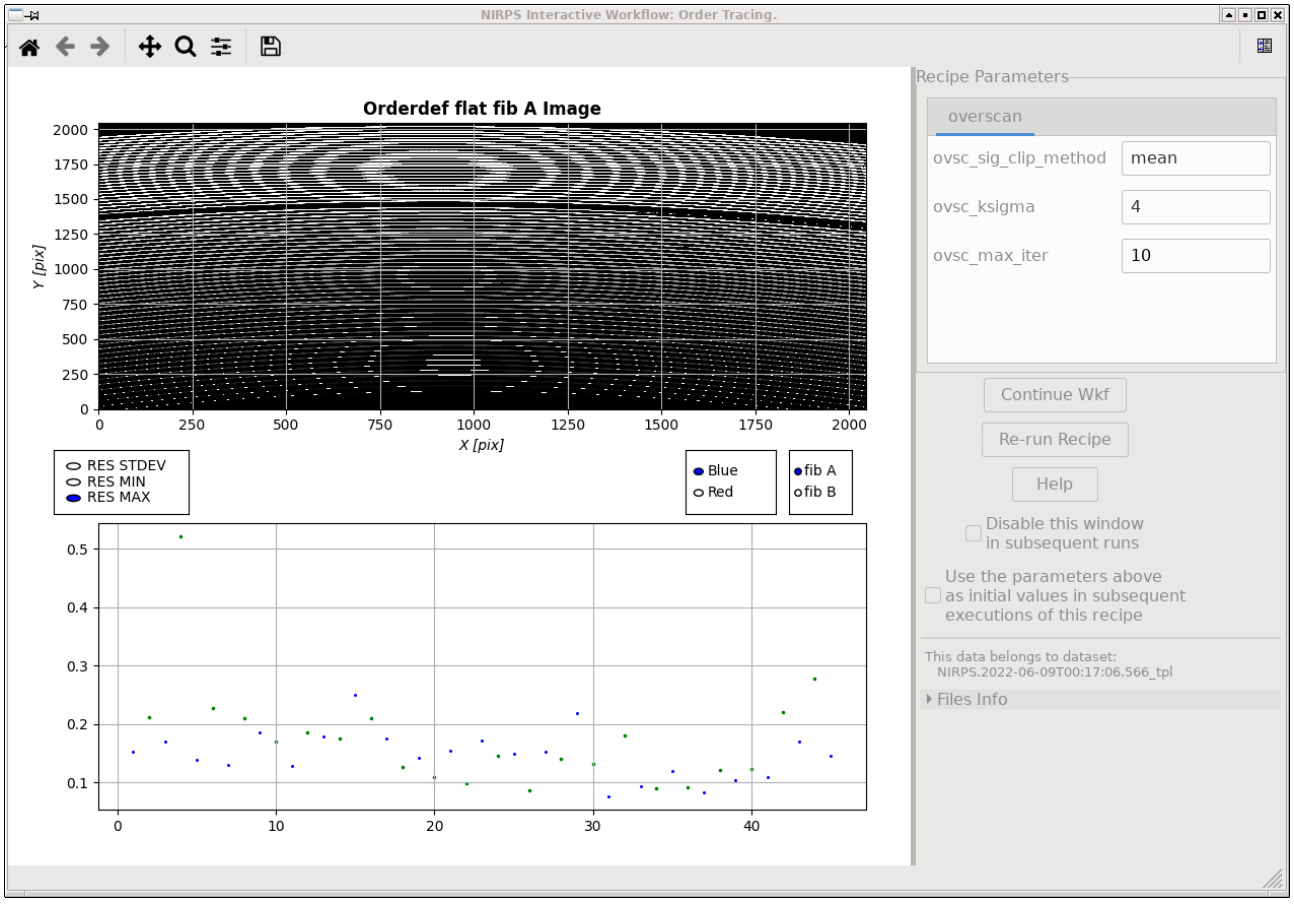

Figure 7.2.1: *Orderdef Interactive workflow.*

7.3 Master Flat

By default the master flat interactive workflow shows at the top left the image corresponding to the stacking of the extracted flat field spectra from both detectors, end to end (FS2D format) for fibre A (Fig. 7.3.1). An FS2D image is a two-dimensional extracted spectrum image in which the one-dimensional spectrum of each order is

contained by in a horizontal line. The wavelength runs along the horizontal axis, increasing from left to right. The size of the FS2D image is chosen to fully contain the extracted spectra.

The plot of the spectrum for a given order (by default number 10; this can be changed clicking on the horizontal slider labeled “order no” at the top of the GUI) is displayed below on the middle-level left panel. The SNR as a function of the order number is plotted in the bottom left panel.

The user may choose to display also results for other products (`FS2D_B`, `FSPECTRUM_A/B`, `BLAZE_A/B`, `ORDER_PROFILE_A/B`) clicking on the respective buttons form the menu on the top left of the GUI.

The interactive workflow allows the user to change the parameters (top right area) controlling the recipe behaviour: determination of the overscan (ignored for NIRPS), the frame stacking, the inter-order background determination, and the spectrum extraction.

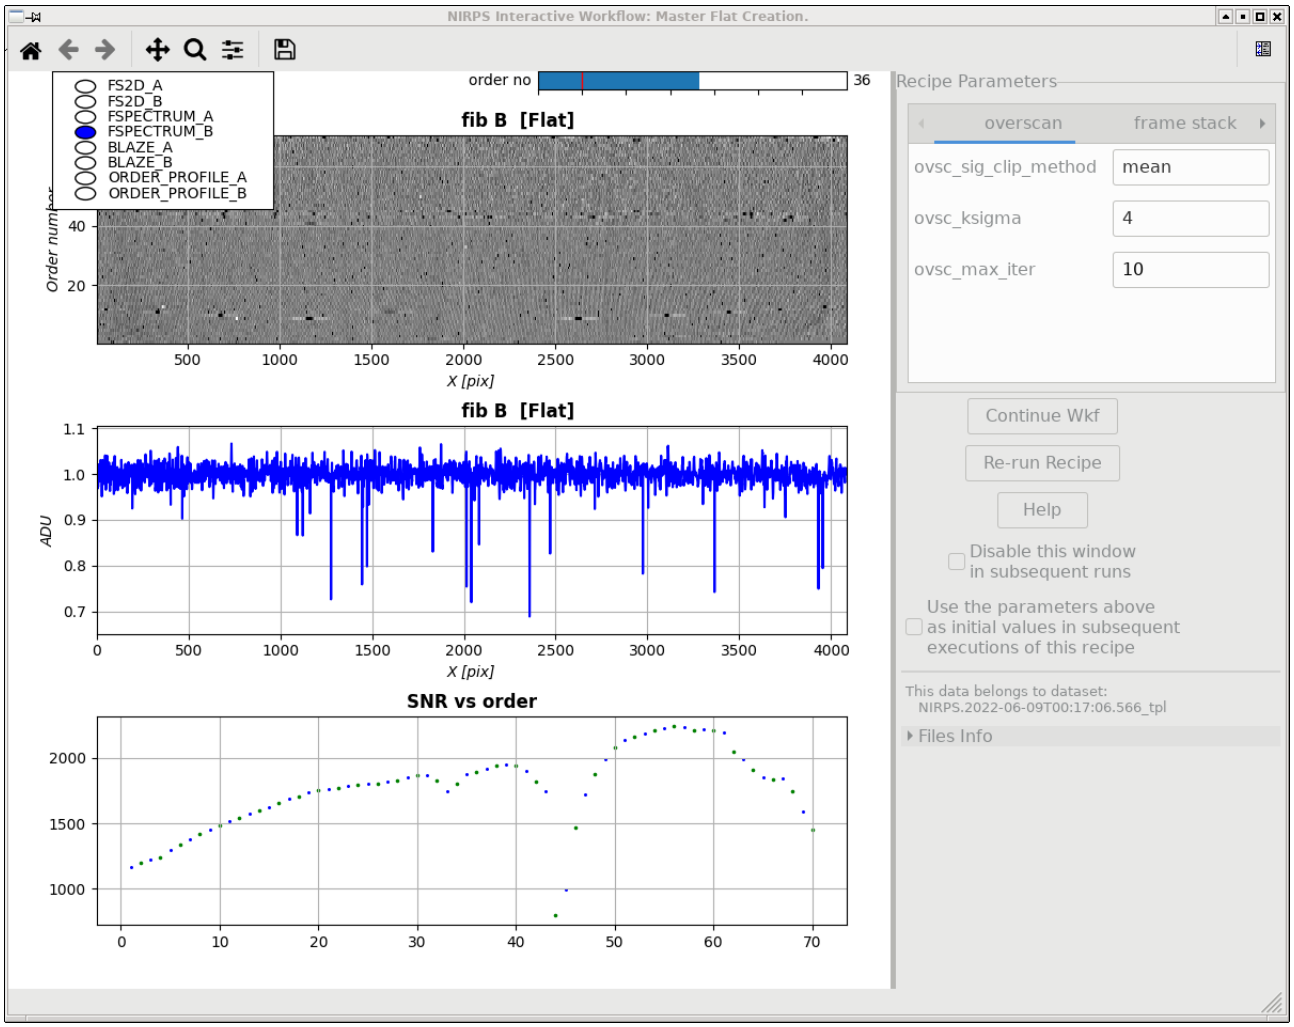

Figure 7.3.1: *Master Flat Interactive workflow.*

7.4 Wavelength Calibration: processing of frames where both fibres are illuminated by a Fabry Perot

This interactive workflow shows on the top left panel the stacking of the extracted Fabry Perot spectra, end to end (S2D format; Fig. 7.4.1). The user may display the image for fibres A or B, or the stacking of the derived blaze functions for fibres A or B (they can be selected by pressing the buttons on the menu in the top left corner).

The middle left panel shows the extracted spectrum of the order number indicated by the horizontal slider (labeled “order no”; default is 10). At the bottom left panel the GUI shows the scatter plot of the number of detected lines for the same order, that is displayed on the middle left panel

The user may change the recipe parameters from the area on the right side of the GUI.

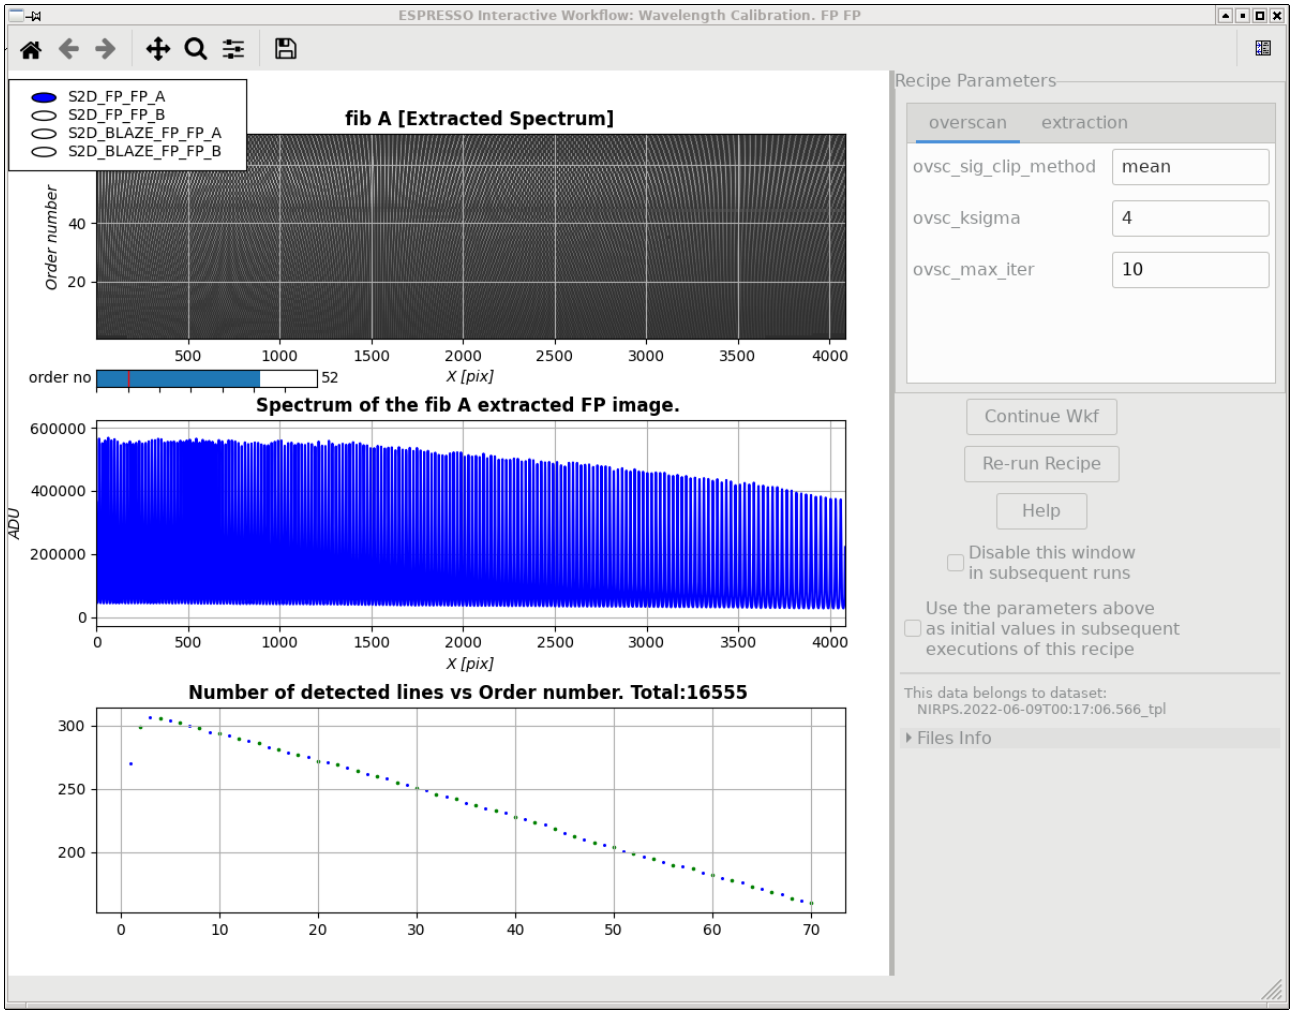

Figure 7.4.1: *Wavelength Calibration (FP,FP) Interactive workflow.*

7.5 Wavelength Calibration: processing of FP,THAR or THAR,FP data

This interactive workflows shows at the top left panel an image of the stacked to 2D format extracted Fabry Perot an ThAr 1D spectra, end to end (S2D format, see Fig. 7.5.1) for THAR,FP. By default this is for fibre A, but the users may display similar images for fibre B (FP spectrum) or for the extracted blaze spectrum for fibres A or B. This is controlled pressing the buttons in the menu labeled “Control top panels” at the lower left of the GUI.

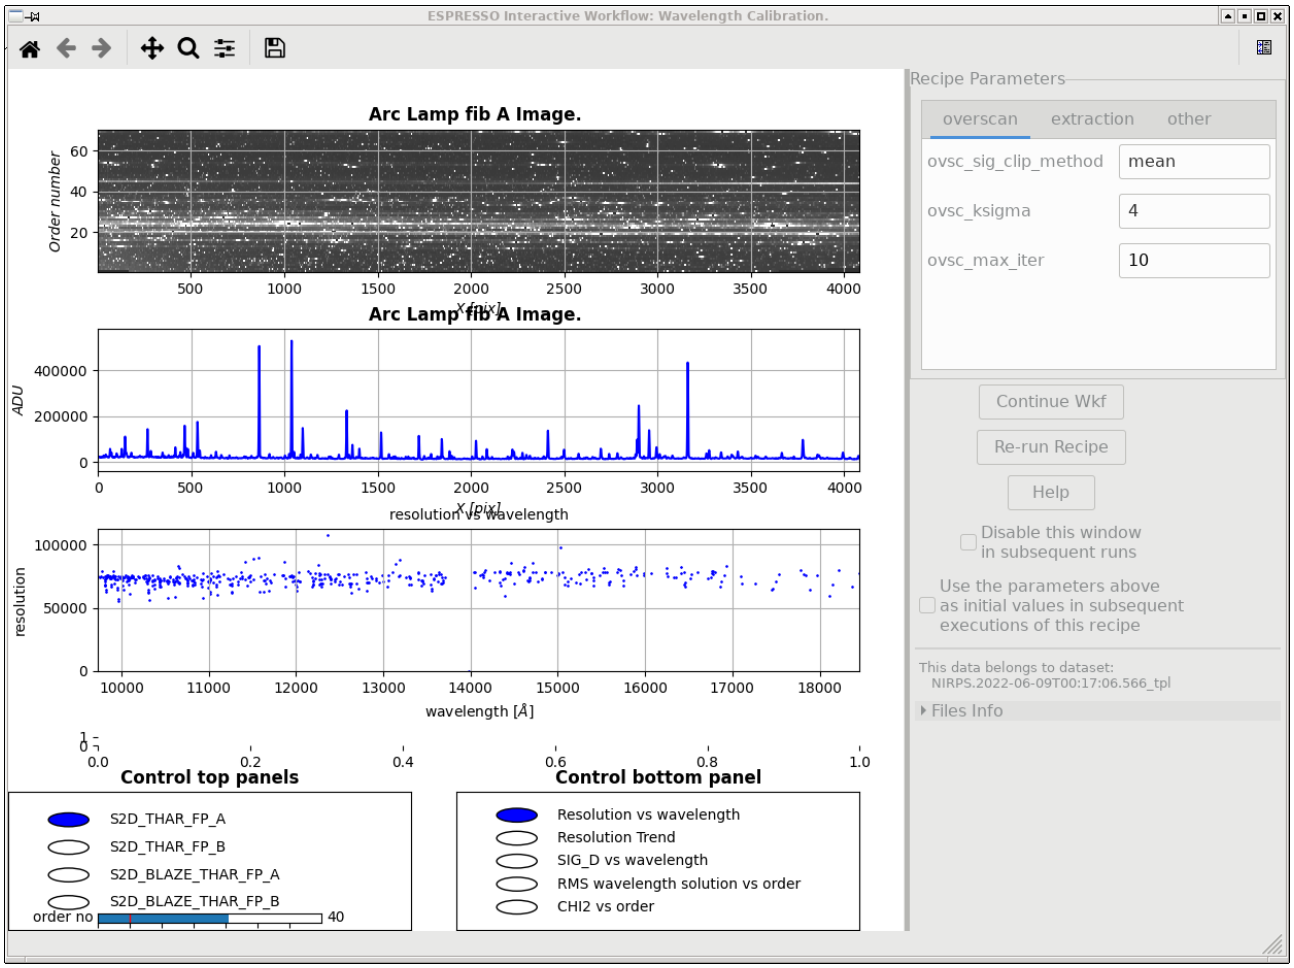

Figure 7.5.1: *Wavelength Calibration (THAR,FP) Interactive workflow.*

In the second (from top to bottom) left panel the GUI shows the extracted 1D spectrum corresponding to the image at the top and to the order number shown in the horizontal slider (on the bottom of the GUI, labeled with “order no”).

The third panel on the left shows the scatter plot of the resolution versus wavelength by default. The user may change this to display the resolution trend in different detector positions (this is shown with a different color scale) or the effective Fabry Perot cavity length (D) measured by the pipeline or its sigma, versus wavelength (LL), or the RMS of the wavelength solution versus the order number, or the computed CHI2 versus the order

number. This is controlled pressing the buttons in the menu labeled “Control bottom panel” at the lower center of the GUI.

The user may change the recipe parameters that control the extraction, the polynomial fit of D, and the polynomial wavelength solution; the parameters for the overscan detection are ignored for NIRPS, but are left for consistency with other workflows.

The interactive GUI showing results of FP,THAR reduction is similar (Fig. 7.5.2).

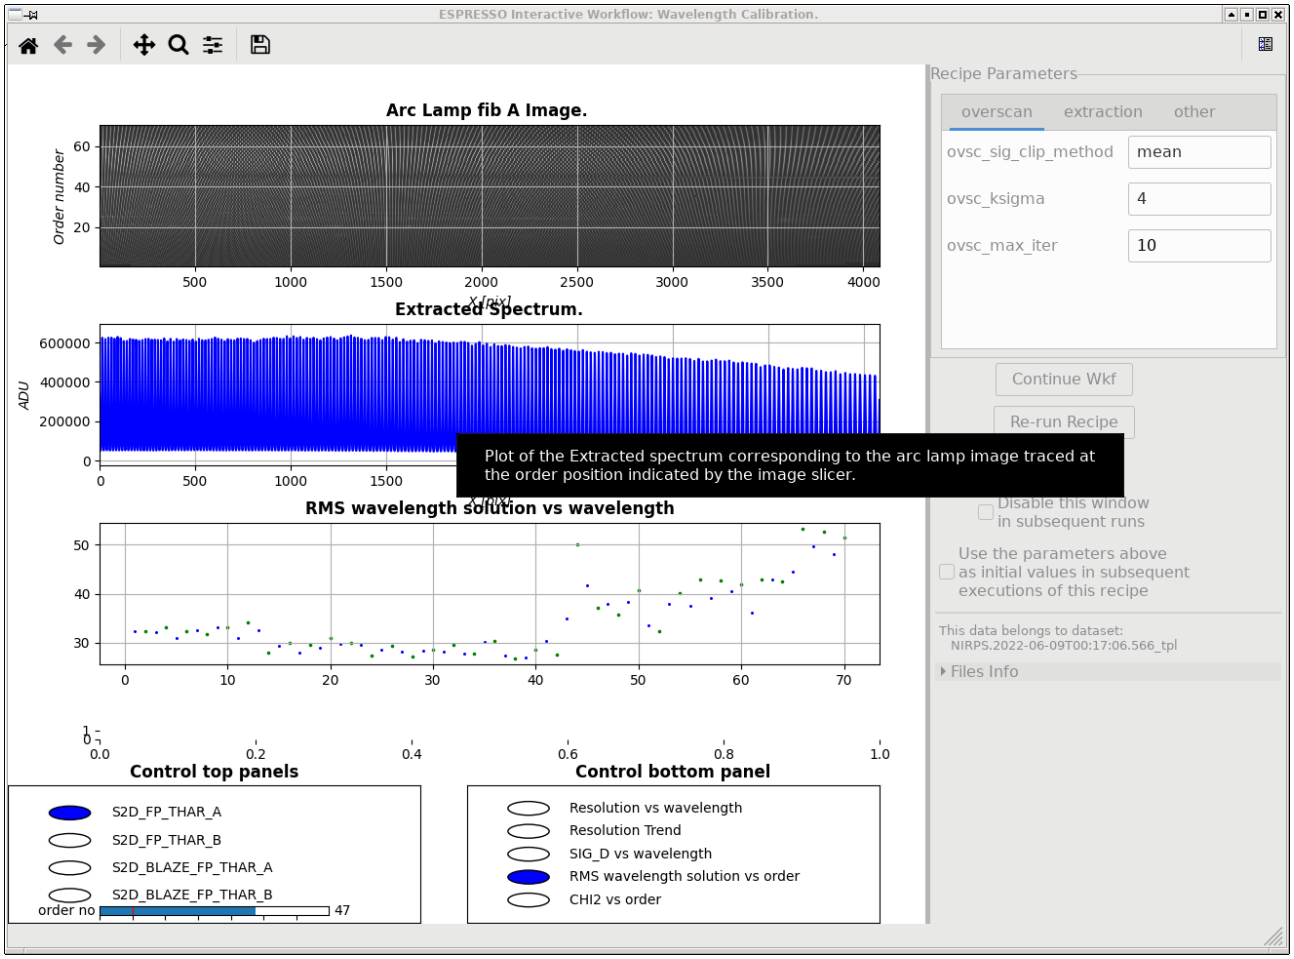

Figure 7.5.2: *Wavelength Calibration (FP,THAR) Interactive workflow.*

Note: If the peaks of the emission lies are not enough intense, the recipe can miss some identifications. In the unfortunate event that none of the peaks that should be identified in both orders are detected, the recipe cannot "connect" the wavelenght solution between orders and it fails. The error message in this case is:

```
[ ERROR ] espdr_wave_THAR: espdr_get_all_FP_ll_whole_image failed: Input data do not match
```

In some cases, this situation can be avoided by changing the recipe parameter `fp_extraction_ksigma` in the interactive window (tab: extraction).

7.6 Cross fibre contamination determination

This interactive workflow shows at the top left panel the image of the spectra for fibre A, used to compute the contamination of fibre B on fibre A (Fig. 7.6.1). The user may alternatively display the corresponding spectrum for fibre B or the image describing the contamination of the Fabry Perot image. The middle canvas shows the extracted spectrum corresponding to the order given in the horizontal slider. The canvas at the bottom shows the scatter plot of the maximum flux as a function of the order number.

The user may change parameters controlling the overscan determination, the inter-order background computation, and the extraction. The user may change the parameters that control the overscan, the inter-order background determination, the extraction and the polynomial fit degree used to fit the efficiency.

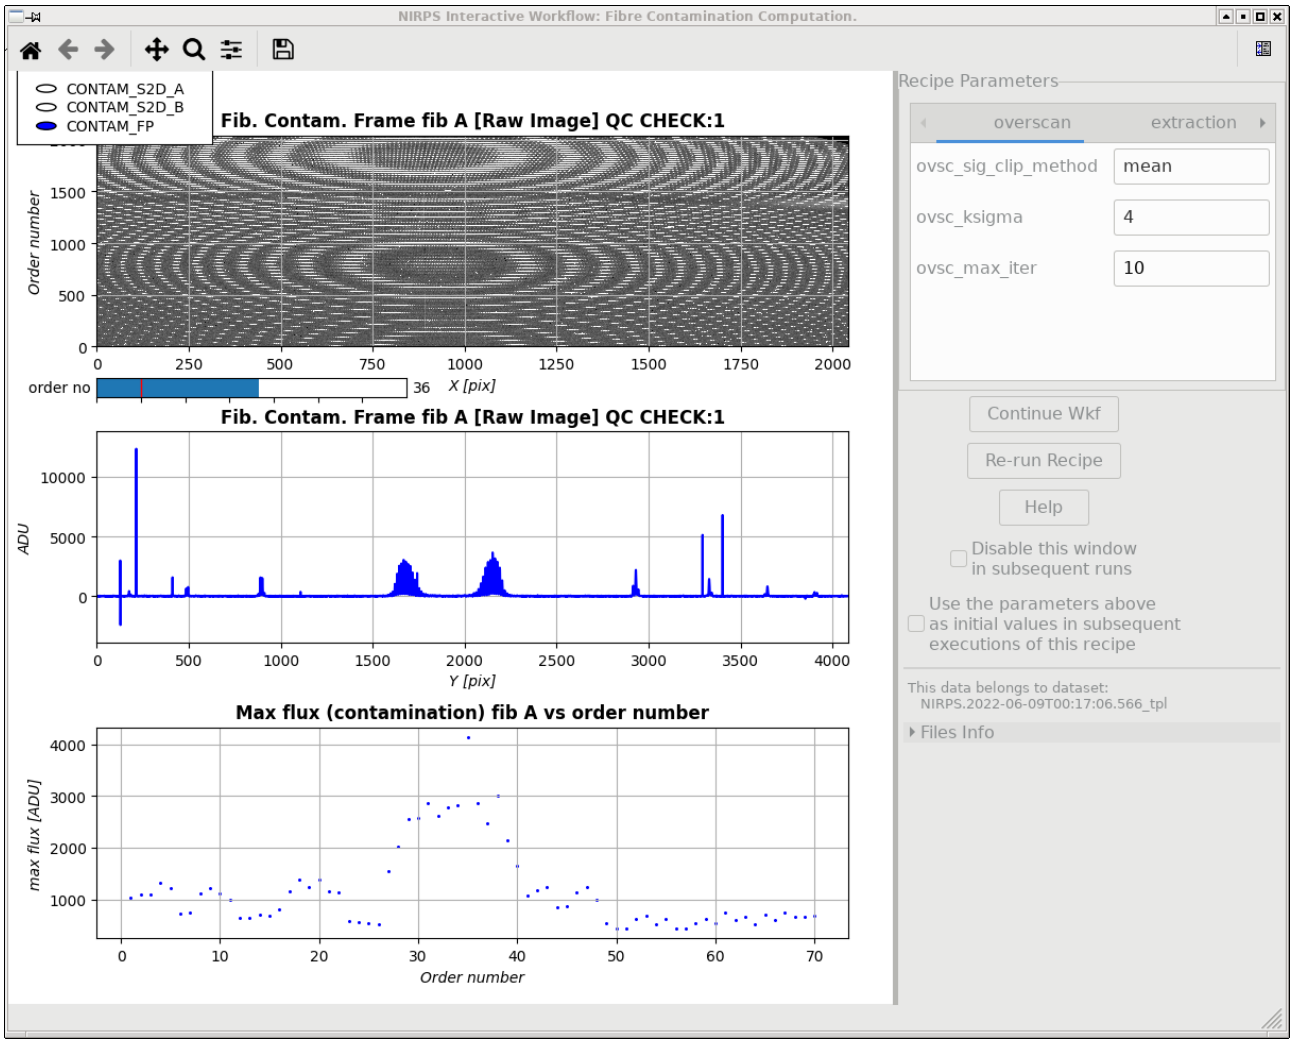

Figure 7.6.1: *Contamination data Interactive workflow.*

7.7 Relative Fibre efficiency determination

This interactive GUI shows at the top left an image in S2D format of the extracted spectra ratio B-to-A (Fig. 7.7.1). In the middle left panel it shows the relative efficiency as function of order number and at the bottom left – the SNR for each order.

The user may change the parameters that control the overscan, the inter-order background determination, the extraction and the polynomial fit degree used to fit the efficiency.

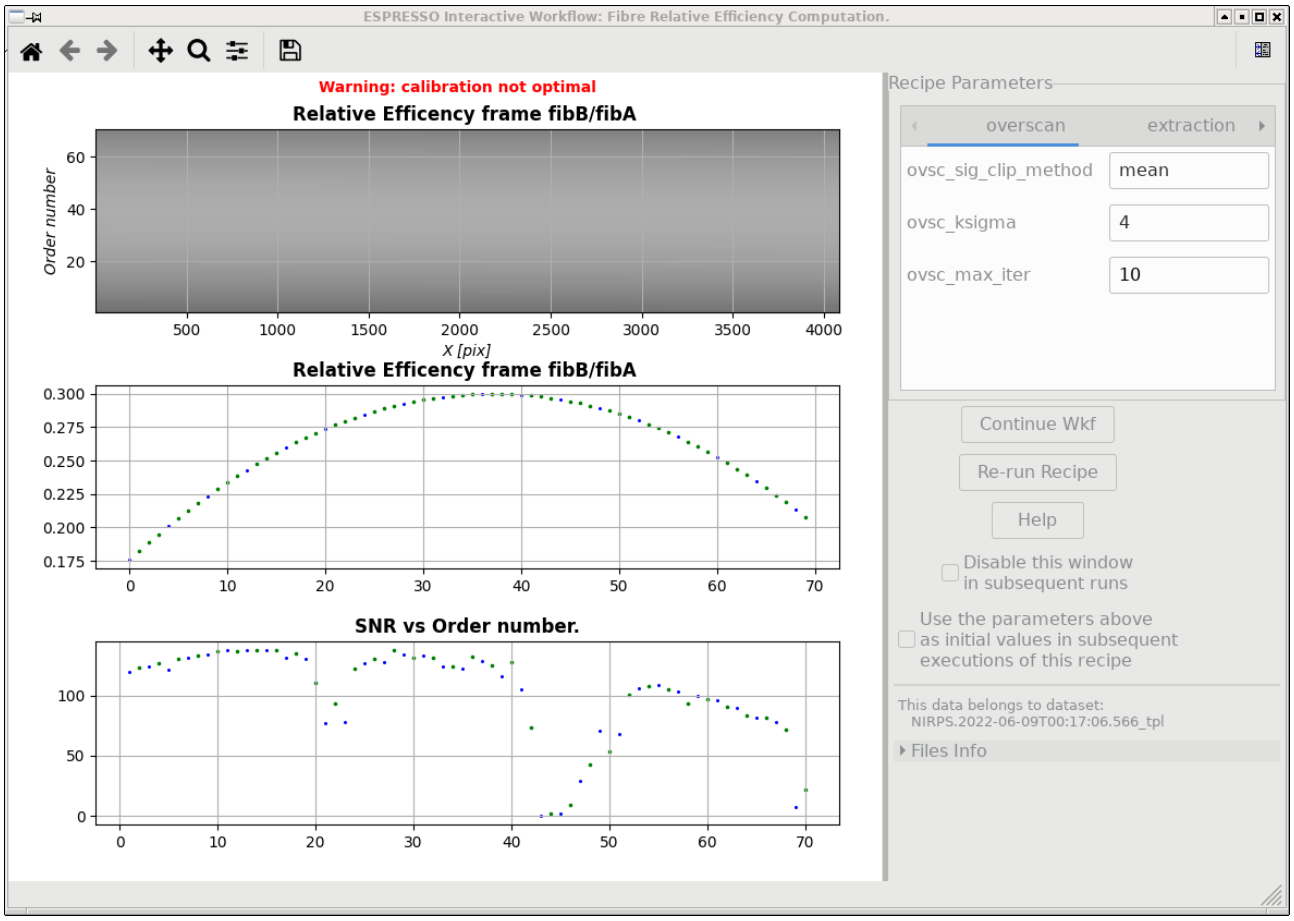

Figure 7.7.1: *Relative fibre efficiency determination Interactive workflow.*

7.8 Flux calibration

This interactive GUI shows at the top left an image in S2D format of the combined detectors extracted flux STD star spectra for fibre A (Fig. 7.8.1). Alternatively, the user may direct the GUI to display the image of the blaze function in S2D format by pressing the corresponding button on the menu (labeled “Control upper panels”) at the bottom left of the GUI.

In the middle left panel shows the extracted spectrum for the order corresponding to what the horizontal slider displays (at the very bottom, labeled “order no”; default is 10).

The third (from the top) left panel shows the computed efficiency over the entire wavelength range. Alternatively, the user may direct it display one of the following:

– the computed efficiency before the polynomial fit,

– the spectrum of the average flux observed for the flux STD star,

– the corresponding S1D extracted spectrum,

– the corresponding S1D measured energy spectrum.

by pressing the corresponding button on the menu (labeled “Control bottom panel”) at the bottom of the GUI.

The user may change the recipe parameters that control the overscan (ignored for NIRPS), the inter-order background determination, the extraction and the polynomial fit degree used to fit the efficiency.

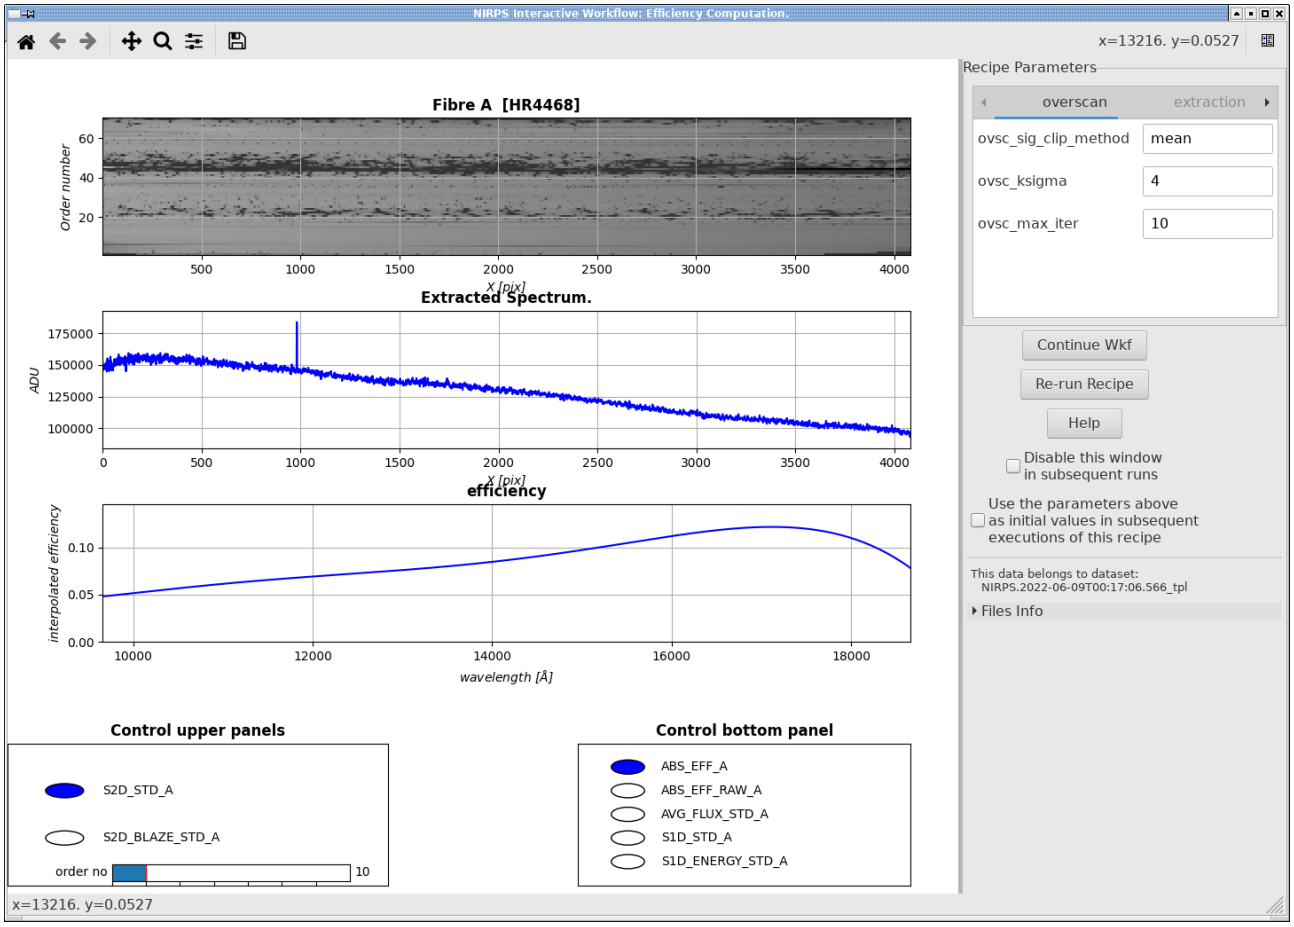

Figure 7.8.1: *Flux calibration determination Interactive workflow.*

7.9 Science Data Reduction

This interactive GUI shows at the top left an image in S2D format of the combined detectors extracted object spectra for fibre A (Fig. 7.9.2). Alternatively, the user may display the image of the spectrum corresponding to fibre B in S2D format, or and image of the Cross Correlation Function (CCF).

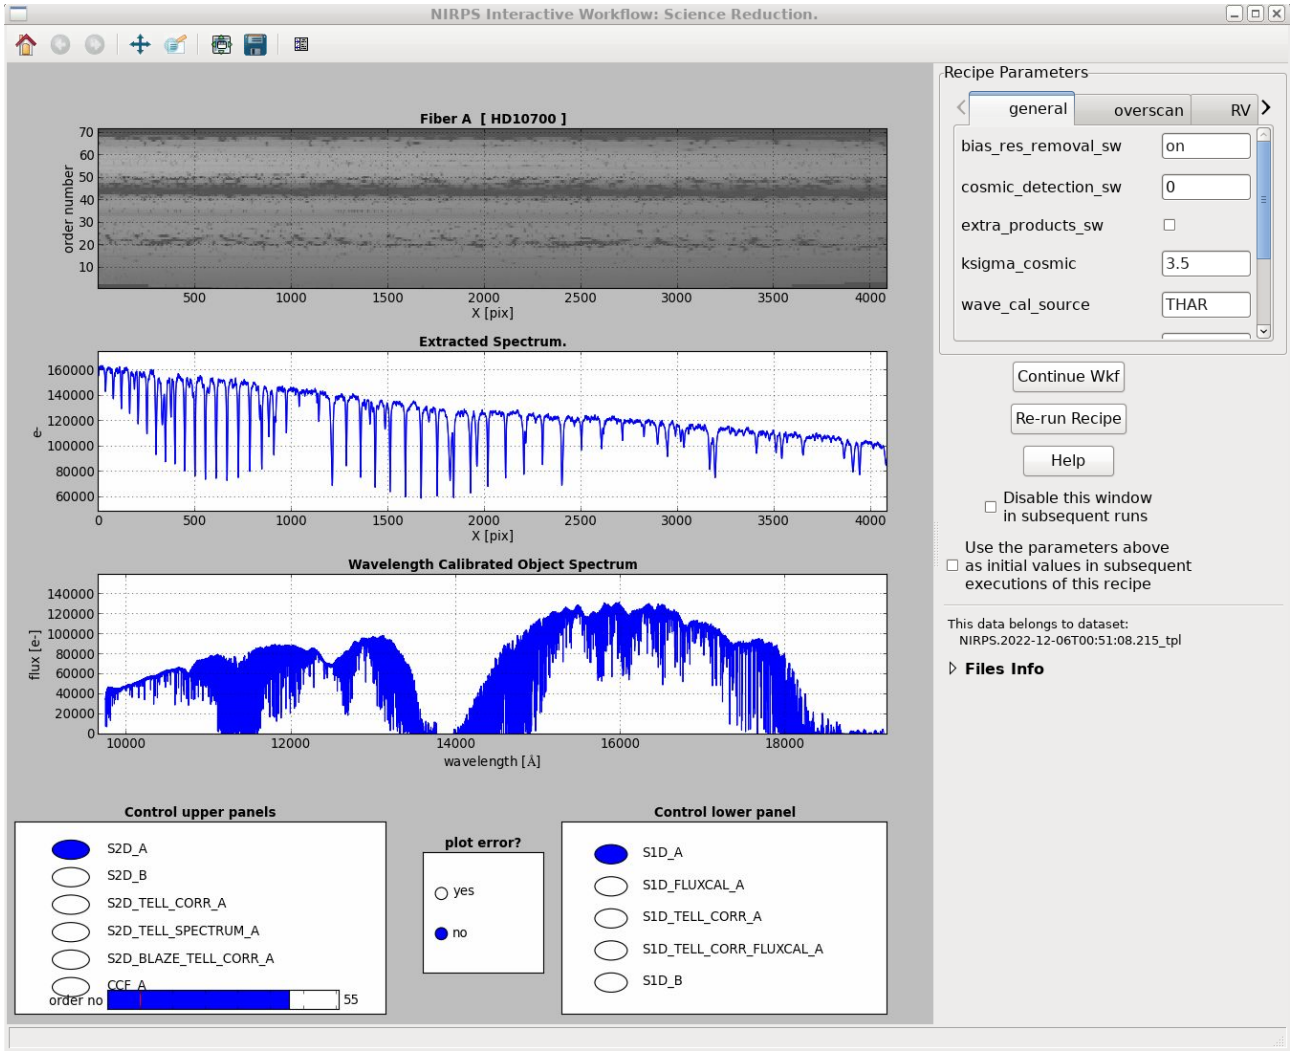

Figure 7.9.1: *Science data reduction Interactive workflow. The spectrum without telluric correction is shown on the third panel on the left.* .

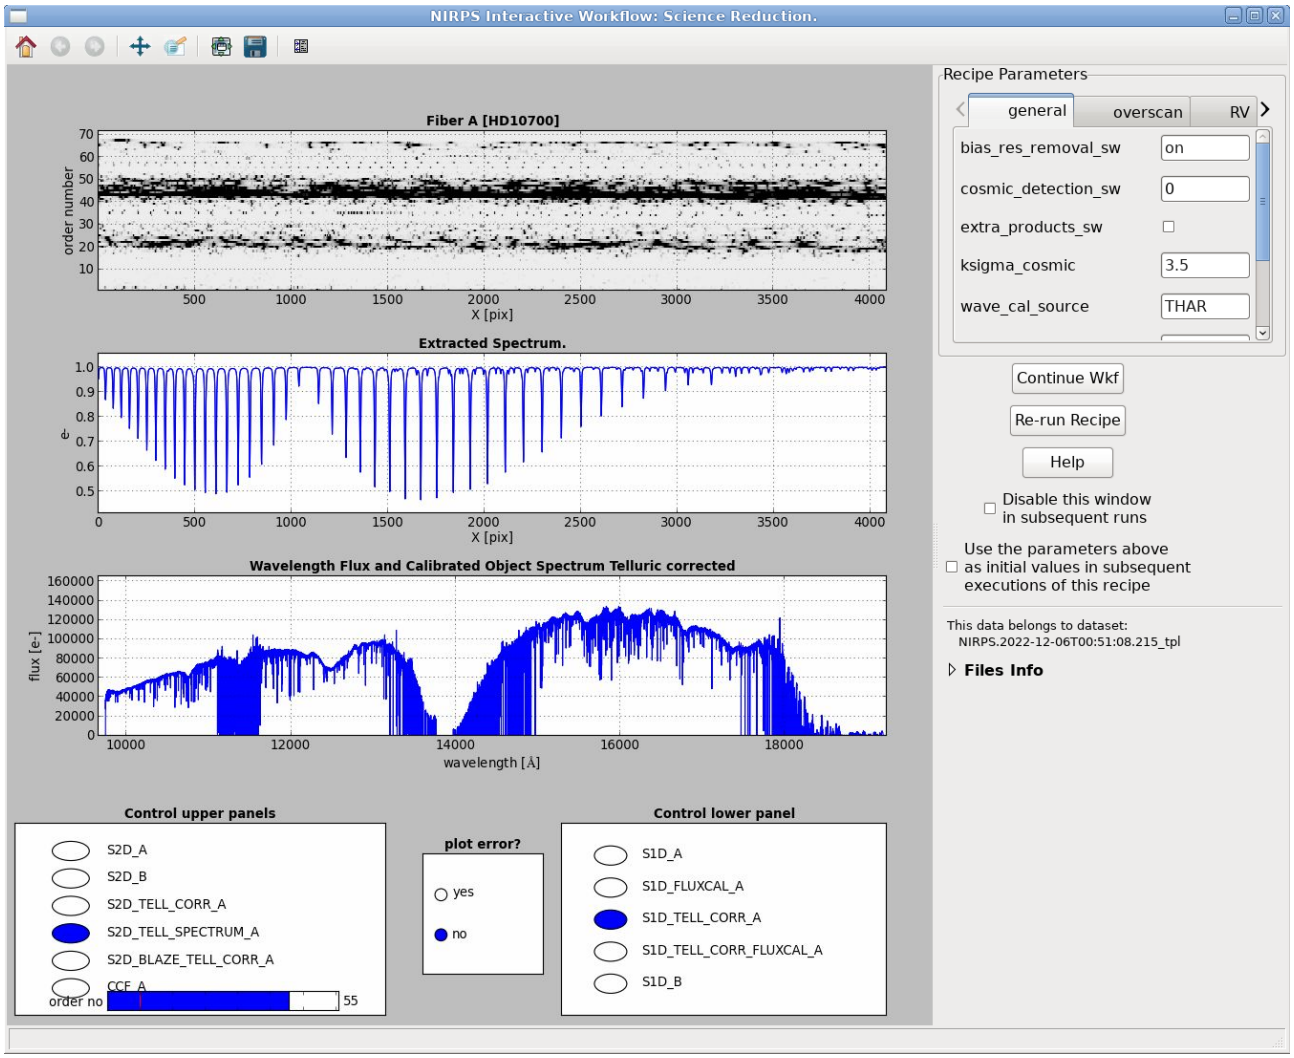

Figure 7.9.2: *Science data reduction Interactive workflow. The spectrum with telluric correction is shown on the third panel on the left. The spectum on the second panel for the second plot is the atmospheric transmission, normalized to unity..*

In the middle left panel the extracted spectrum is shown for the order number corresponding to what is set at the horizontal slider (labeled “order no”). This panel can display also one of the following:

– the spectrum of the object (default),

– the flux calibrated spectrum of the object,

– the spectrum of what is in the fibre B.

The control over the first and the second (from the top) panels is exercised from the menu labeled “Control upper panels” at the bottom left of the GUI. The user can also plot the error-bar (time consuming) on each S1D spectra (see the menu “plot error?” at the bottom of the GUI).

The third left panel from the top can display:

– sky corrected and wavelength calibrated object spectrum in fibre A (default option)

– sky corrected, wavelength calibrated and flux calibrated object spectrum in fibre A

– wavelength calibrated object spectrum in fibre A

– wavelength calibrated and flux calibrated object spectrum in fibre A

– IDP spectrum of what is in fibre B.

The control over the third panel is exercised from the menu labeled "Control lower panel" at the bottom of the GUI.

Note: for extragalactic targets, identified by a non-null value of the header keyword `HIERARCH ESO OCS OBJ Z EM`, the cross correlation function is not calculated.

Useful tips:

The user may change the parameters that control the inter-order background determination, the extraction, the radial velocity computation, the rejection of cosmics, the drift correction.

For very low S/N, readout limited regime or close to it: it is recommended to use the smooth option for sky subtraction, reducing the readout noise contribution from the sky subtraction step.

The recipe parameter **background_sw** removes the intra-order diffuse background contamination (it must not be misinterpreted as sky background removal, which is regulated by **sky_sub_method** and is active only for OBJECT,SKY observations). The intra-order background is evaluated on boxes of size **sci_bkgr_grid_size_x** $\times$ **sci_bkgr_grid_size_y** and then interpolated all over the frame. The diffuse background is proportional to the dispersed light inside the spectrograph (and thus, to a good approximation, to the total flux recorded in the detector). If the diffuse background is not measured correctly (as it happens when it is negligible), correcting for it will only contribute additional noise to the flux. The FP etalon contributes significantly to the inter-order background and for exposures with FP on fiber B the inter-order background correction flag is set by default to ON. In most cases, stellar light does not contribute significantly to inter-order background and for exposures with SKY on fiber B the inter-order background correction flag is set by default to OFF. If any exposure points to the SKY on fiber B and nevertheless has a very high target flux in fiber A, the user should change the inter-order correction flag to ON.

7.9.1 Telluric Correction

The NIPRS pipeline can apply a simple telluric model to corret for the absorption of the Earth atmosphere. Water is the main absorpber, but other species such as carbon dioxide, oxygen and methane contribute significantly at certain wavelengths. This step is applied to the data by default, although a version of the spectrum without the correction is also available to the users. The model is a line-by-line radiative transfer code assuming a single atmospheric layer, based on the `HITRAN` (high-resolution transmission molecular absorption) database line list. The atmospheric properties are determined from a fit to selected telluric lines. For furthrer details see Allert et al (2022).

7.9.2 Cosmic rays cleaning

The science recipe has two algorithms for cosmic ray cleaning. They user can decide to use none, one or both by setting the corresponding recipe parameter in the interactive window.

The first option is the LA Cosmic algorithm (van Dokkum 2001, PASP 113, 1420) to flag cosmic rays in the raw frame. This can be enabled or disabled by setting the recipe parameter `cosmic_detection_sw` to 1 or 0, respectively. This parameter is located in the tab "General" of the interactive window. Flagged pixels are ignored during the extraction of the 1-dimensional spectrum. The tab "LACOSMIC" in the interactive window allows the user to specify the recipe parameters that control the algorithm itself. By default, the workflow switches on/off the algorithm and loads the optimal recipe parameters for a given observing mode. We refer the user to the pipeline manual for more information on the recipe parameters. The user might want to inspect the products with category `CRH_MAP` and `CCD_CORR_SCIENCE` to assess the quality of the flagging process. These products are available in the reflex temporary products directory and are produced only if the recipe parameter in the "General" tab `extra_products_sw` is up. Future versions of the workflow will allow inspection on the interactive window.

The second algorithm is a "k-sigma' clipping"; it operates on the extracted spectrum and flags pixels that deviate more than a certain factor (specified by the recipe parameter `ksigma_cosmic` times the noise from the model profile. This parameter is located in the tab "General" of the interactive window. Set `ksigma_cosmic` to a negative number to switch off this algorithm.

The general recommendation is to turn on both algorithms, but pay attention whether very intense targets or very intense exposures of the Fabry Perot fibre (which have very sharp edges in the spectral profile, in particular with high binning modes), do generate too many false positive detections. In this case, it is recommended either to switch the LA Cosmic algorithm off, or to increase the `f_lim sigma_lim` thresholds. If a cosmic ray is only partially detected, one might dilate the cosmic-ray mask by setting `post-filter-x` and `post-filter-y` parameters to 1 or 2 (higher values are not recommended).

The user has the option to provide a mask with cosmic rays detection (category `CRH_MAP`). If provided, it will switch off the execution of the LA Cosmic algorithm. In order to be used by the workflow and the recipes, the user-provided mask has to:

– be located in the `RAWDATA_DIR` with other input files

– have the category `CRH_MAP` (header keyword `HIERARCH ESO PRO CATG`)

– have the same format of `CRH_MAP` as the mask generated by the LA Cosmic algorithm in the science recipe `espr_sci_red`; this includes the same dimensions and number of extensions, and the same pixel value convention: cosmic rays are set to 1, unaffected pixels – to 0 (integer values)

– have the header keyword `ARCFILE` of the same value as the science frame it is supposed to be associated to.

It is recommended to run first the reduction with the LA Cosmic algorithm on, with the option `extra_products_sw`, and then edit the `CRH_MAP` produced by the workflow according to the user's needs. It is also possible to run a third-party cosmic ray detection algorithm on `CCD_CORR_SCIENCE` and set the flagged pixel to 1 in `CRH_MAP`.

Note that the inclusion of an user-provided `CRH_MAP` replaces the LA Cosmic algorithm in the science recipe, but the "k-sigma clipping" algorithm can still be executed.

7.9.3 Radial Velocity

The computation of radial velocity is a key-feature in the NIRPS pipeline. One might decide to fine-tune the recipe parameters to improve the calculation or to run a dedicated workflow and recipe to execute only this step. A dedicated workflow is available in the so-called NIRPS-DAS pipeline, named `espda_starIII_wkf` we refer the user to the NIRPS-DAS documentation for this option.

If one wants to optimize the computation of the radial velocity in the NIRPS pipeline, the recipe parameters to optimize are located into the RV tab of the interactive window. The general recommendation is to start with a large range `rv_range` and a course grid `rv_step` to locate the minimum. Then a fine-tuning of the range and grid are necessary to increase the precision of the measurement.

However, M-dwarfs are particularly prone to bad convergence in the CCV computation if the RV range is large. It is not recommended to derive RVs using an RV range larger than 20 km/s for M dwarfs.

When computing the radial velocity of the same target over a certain period of time, it is important to use the same template (`mask_table_id`) to remove systematic effects between measurements at different epochs.

8 Frequently Asked Questions

- **The error window fills the whole screen - how can I get to the `Continue`/`Stop` buttons?**

  Press the Alt key together with your left mouse button to move the window upwards and to the left. At the bottom the `Continue`/`Stop` buttons will be visible. This bug is known but could not yet be fixed.

- **I tried to `Open` (or `Configure`) an `Actor` while the workflow is running and now it does not react any more. What should I do?**

  This is a limitation of the underlying Kepler engine. The only way out is to kill the workflow externally. If you want to change anything while a workflow is running you first need to pause it.

- **After a successful reduction of a data set, I changed this data set in some way (e.g. modified or removed some files, or changed the rules of the Data Organizer). When I restart Reflex, the Data Set Chooser correctly displays my new data set, but marks it as "reduced ok", even though it was never reduced before. What does this mean?**

  The labels in the column "Reduced" of the Data Set Chooser mark each dataset with "OK", "Failed" or "-". These labels indicate whether a data set has previously successfully been reduced at least once, all previous reductions failed, or a reduction has never been tried respectively. Data sets are identified by their name, which is derived from the first science file within the data set. As long as the data set name is preserved (i.e. the first science file in a data set has not changed), the Data Organizer will consider it to be the same data set. The Data Organizer recognizes any previous reductions of data sets it considers to be the same as the current one, and labels the current data set with "OK" if any of them was successful, even if the previously reduced data set differs from the current one.

  Note that the Product Explorer will list all the previous reductions of a particular data set only at the end of the reduction. This list might include successful and/or unsuccessful reduction runs with different parameters, or in your case with different input files. The important fact is that these are all reductions of data sets with the same first raw science file. By browsing through all reductions of a particular raw science file, the users can choose the one they want to use.

- **Where are my intermediate pipeline products?** Intermediate pipeline products are stored in the directory `<TMP_PRODUCTS_DIR>` (defined on the workflow canvas, under Setup Directories) and organised further in directories by pipeline recipe.

- **Can I use different sets of bias frames to calibrate my flat frames and science data?** Yes. In fact this is what is currently implemented in the workflow(s). Each file in a DataSet has a purpose attached to it (Forchì (2012)). It is this purpose that is used by the workflow to send the correct set of bias frames to the recipes for flat frame combination and science frame reduction, which may or may not be the same set of bias frames in each case.

- **Can I run `Reflex` from the command line?** Yes, use the command:

  ```
  esoreflex -n <workflow_path>/<workflow>.xml
  ```

  The -n option will set all the different options for Kepler and the workflows to avoid opening any GUI elements (including pipeline interactive windows).

  It is possible to specify workflow variables (those that appear in the workflow canvas) in the command line. For instance, the raw data directory can be set with this command:

```
esoreflex -n -RAW_DATA_DIR <raw_data_path> \
             <workflow_path>/<workflow>.xml
```

You can see all the command line options with the command `esoreflex -h`.

Note that this mode is not fully supported, and the user should be aware that the path to the workflow must be absolute and even if no GUI elements are shown, it still requires a connection to the window manager.

- **How can I add new actors to an existing workflow?** You can drag and drop the actors in the menu on the left of the `Reflex` canvas. Under `Eso-reflex -> Workflow` you may find all the actors relevant for pipeline workflows, with the exception of the recipe executer. This actor must be manually instantiated using `Tools -> Instantiate Component`. Fill in the "Class name" field with `org.eso.RecipeExecuter` and in the pop-up window choose the required recipe from the pull-down menu. To connect the ports of the actor, click on the source port, holding down the left mouse button, and release the mouse button over the destination port. Please consult the Reflex User Manual (Forchì (2012)) for more information.

- **How can I broadcast a result to different subsequent actors?** If the output port is a multi-port (filled in white), then you may have several relations from the port. However, if the port is a single port (filled in black), then you may use the black diamond from the toolbar. Make a relation from the output port to the diamond. Then make relations from the input ports to the diamond. Please note that you cannot click to start a relation from the diamond itself. Please consult the Reflex User Manual (Forchì (2012)) for more information.

- **How can I manually run the recipes executed by Reflex?** If a user wants to re-run a recipe on the command line he/she has to go to the appropriate reflex_book_keeping directory, which is generally reflex_book_keeping/<workflow>/<recipe_name>_<number> There, subdirectories exist with the time stamp of the recipe execution (e.g. 2013-01-25T12:33:53.926/). If the user wants to re-execute the most recent processing he/she should go to the `latest` directory and then execute the script `cmdline.sh`. Alternatively, to use a customized `esorex` command the user can execute

```
ESOREX_CONFIG="INSTALL_DIR/etc/esorex.rc"
PATH_TO/esorex --recipe-config=<recipe>.rc <recipe> data.sof
```

where INSTALL_DIR is the directory where Reflex and the pipelines were installed.

If a user wants to re-execute on the command line a recipe that used a specific raw frame, the way to find the proper data.sof in the bookkeeping directory is via `grep <raw_file> */data.sof`. Afterwards the procedure is the same as before.

If a recipe is re-executed with the command explained above, the products will appear in the directory from which the recipe is called, and not in the reflex_tmp_products or reflex_end_products directory, and they will not be renamed. This does not happen if you use the `cmdline.sh` script.

- **If I enter "-" into an empty integer parameter of an interactive window it is automatically completed to "-1". Why?**

  The parameters are validated for correctness according to their type (e.g. string, integer, float). In the case of an integer or float parameter "-" alone is considered an invalid input and is therefore automatically completed to "-1". This is part of the validation of input done by the WxPython library.

- **Can I reuse the bookkeeping directory created by previous versions of the pipeline?**

  In general no. In principle, it could be reused if no major changes were made to the pipeline. However there are situations in which a previously created bookkeeping directory will cause problems due to pipeline versions incompatibility. This is especially true if the parameters of the pipeline recipes have changed. In that case, please remove the bookkeeping directory completely.

- **How to insert negative values into a textbox?**

  Due to a bug in wxPython, the GUI might appear to freeze when attempting to enter a negative number in a parameter's value textbox. This can be worked around by navigating away to a different control in the GUI with a mouse click, and then navigating back to the original textbox. Once focus is back on the original textbox the contents should be selected and it should be possible to replace it with a valid value, by typing it in and pressing the enter key.

- **I've updated my Reflex installation and when I run esoreflex the process aborts. How can I fix this problem?**

  As indicated in Section 3, in case of major or minor (affecting the first two digit numbers) Reflex upgrades, the user should erase the `$HOME/KeplerData`, `$HOME/.kepler` directories if present, to prevent possible aborts (i.e. a hard crash) of the esoreflex process.

- **How can include my analysis scripts and algorithms into the workflow?**

  EsoReflex is capable of executing any user-provided script, if properly interfaced. The most convenient way to do it is through the Python actor. Please consult the tutorial on how to insert Python scripts into a workflow available here: `www.eso.org/sci/data-processing/Python_and_esoreflex.pdf`

9 Troubleshooting

In this section we describe some of the problems that may occur when reducing the NIRPS data with Reflex. For a more comprehensive description we refer the user to the NIRPS pipeline user manual (`http://www.eso.org/sci/software/pipelines/`). The user may also want to consult the Data Reduction F.A.Q. page at www.eso.org/sci/data-processing/faq.html.

1. **I downloaded the data from the ESO archive, put them into a new directory, tried to run `Reflex` on them, but**

   (a) **The reflex execution fails at some stage.**
   This may happen if one of the files was downloaded only partially (check for a file with the extension `fits.Z.part`. You will have to download that file again in order to have an not corrupted file (and remove the partial one).

   (b) **it fails with error message ":No DataSets have been created, check the data set and the OCA rules." (see Figure ??).**
   This error may be due to the fact that the data provided by the ESO archive are compressed (`<filename>.fits.Z`). Please remember to uncompress the data before executing `Reflex`.

   (c) **all DataSets are greyed out in the DataSets interactive window.**
   The ESO archive used with CalSelector does not always supply all static calibration files. As a consequence some/all DataSets are greyed out because they were missing such required data.
   Missing static calibration should be found by `reflex` in `<install_directory>/calib/<pipeline_version>/cal`.

2. **The "Select DataSets" window displays my DataSets, but some/all of them are greyed out. What is going on?**

   If a DataSet in the "Select DataSets" window is greyed out, then it means that the DataSet that was constructed is missing some key calibration(s) (i.e. the DataSet is incomplete). To find out what calibration(s) are missing from a greyed out DataSet, click on the DataSet in question to highlight it in blue, and then click on the button `Inspect Highlighted`. The "Select Frames" window that appears will report the category of the calibration products that are missing (e.g. `ORDER_PROFILE_A_WKF`, `ORDER_PROFILE_B_WKF`, `FSPECTRUM_A_WKF`, `FSPECTRUM_B_WKF`, `BLAZE_A_WKF`, `BLAZE_B_WKF`). From this the user has then to determine the missing raw data (in this case flat frames for the given instrument mode and detector binning setting). If static calibrations are missing the mechanism unfortunately does not work, but should be found by `reflex` in `<install_directory>/calib/<pipeline_version>/cal`

   A faster way to identify which kind of data are missing is moving the mouse over the grey DataSet (see Fig. 9.0.1).

3. **I'm reducing with ESOReflex NIRPS observations taken between April 1st and the end of July 2023, and a data set is greyed out (incomplete). If I point the cursor on the data set a message appears:**

   - **Missing REL_EFF_B, and**

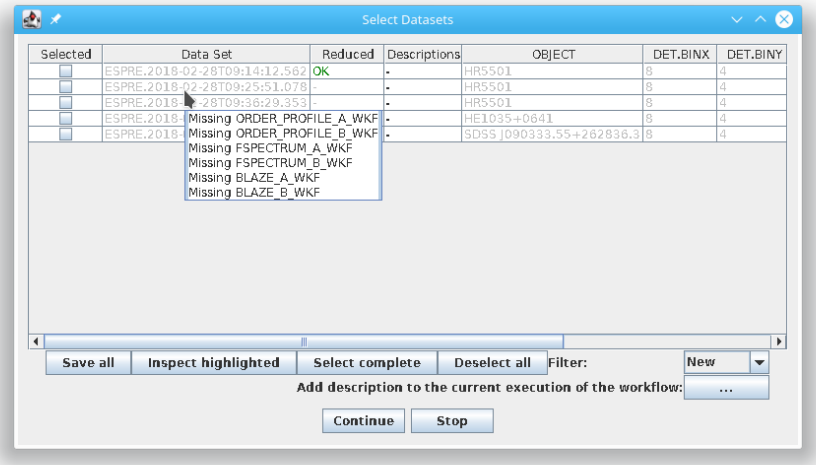

Figure 9.0.1: *This shows how information about missing files in an incomplete DataSet is displayed when moving the mouse over it.*

- **Not enough associated files for organisation rule # 0: requested 5, found 2.**

The former message "Missing REL_EFF_B", means that the raw data required to created the product with PRO.CATG REL_EFF_B are missing.

The latter message "Not enough associated files for organization rule # 0: requested 5, found 2" means that during the file classification and organisation ESOReflex detects that only 2 raw darks are found, less than the required minimum of 5.

These two problems appear during the first four months of operations, because these frames have not been acquired: obtaining the twilight sky frames (DPR.TYPE="EFF,SKY,SKY") that are necessary to determine the relative fibre efficiency is a complex task that was not done at the start of NIRPS operations.

The NIRPS detector is affected by persistence. Therefore, the acquisition of good quality darks (not affected by the persistence) requires to wait a few hours, keeping the instrument idle. During the first months of operations this was not possible because of daily Sun observations; this became possible in Aug.

If the users encounter this problem, they should execute the following commands:

- `cd <path>`
- wget https://ftp.eso.org/pub/dfs/pipelines/nirps/nirps_extra_raw_calibs.tar.gz .
- `tar xvfz nirps_extra_raw_calibs.tar.gz`
- `cd $RAW_DATA_DIR`
- `ln -s <path>/nirps_extra_raw_calibs .`
- restart the workflow.

4. **I downloaded the data from the ESO archive and put them into a new directory, tried to run `Reflex` on them, but some DataSets are greyed out in the DataSets interactive window. How can I be sure to have all the data I need to run successfully the workflow on a given data set?**

   You may miss some calibrations or observations. A useful command to check what input raw types you have are the following (just three of examples):

   - command to check if you have frames to compute the relative fibre efficiency in mode HE:

     ```
     dfits data/*.fits | fitsort DPR.TYPE INS.MODE | grep "EFF"| grep HE
     ```

   - command to check if you have five dark frames stars of mode HE:

     ```
     dfits data/*.fits | fitsort DPR.TYPE INS.MODE TPL.ID TPL.NEXP | grep "DARK"| grep cal_dark | grep HE
     ```

   - command to check if you have led flat frames of mode HE to determine the bad pixel mask:

     ```
     dfits data/*.fits | fitsort DPR.TYPE INS.MODE | grep "LED"| grep HE
     ```

   Similar commands can be given to check for other data types.

5. **Should I change the CALIB_DATA_DIR configuration?**

   This directory is setup automatically to point to the calibration database provided with the pipeline and in principle shouldn't be changed. However, if the static calibration data are present in the RAWDATA_DIR (e.g. calibrations are downloaded from the archive, or copied from ESO-DVD distribution)[5], then you have to set the value of the global variable RAWDATA_DIR to this directory (otherwise an obsolete static calibration file may be selected instead of the most appropriate one).

6. **The data reduction fails at bias or dark reduction step. What can I do?** Please make sure you have at least five input bias and five input dark frames.

7. **The data reduction fails at led-ff reduction step. What can I do?** Please make sure you have at least five input led-ff frames for each exposure time. You need at least two set of frames with different exposure time.

8. **The data reduction fails at the flat reduction step. What I can do?** Please report this problem indicating the full input data set ARCFILE list. You probably have a problematic order determination frames. Usually changing the raw orderdef frames (or removing from the input directory the problematic ones, if you have alternative ones) and restarting the full data reduction solves the problem.

   Another possible source of flat recipe failure is to have in input less than five input flats for each fibre, or a different number of input flats for each fibre. Usual input setting expects ten input raw flats for fibre A and ten for fiber B. This could be due to an incomplete template. You may decide to remove the problematic input flats of that template (the frames with common TPL.START) and use flat frames of the same setting (INS.MODE and DET.BINX) acquired properly by another successful template execution.

9. **"Why the recipe parameters that I specify to "Use the parameters above as initial values in subsequent executions of this recipe" in the interactive window are not used by the recipe?** This can happen if the dataset has been already reduced in the past. The standard behavior for a recipe in interactive actors is to re-use the parameters from the last execution, that are stored in a database. This can happen if the dataset has been already reduced in the past. The standard behaviour for a recipe with interactive

[5] note that the data distributed over DVDs are also available via the ESO science archive, with calibrations selected via the CalSelector tool)

options is to re-use the parameters from the last execution, that are stored in a databasa. This behavior can be changed by editing the configuration of the RecipeLooper actor of each interactive actor and change the "Get initial params from DB" from 'true' to 'false'. Then, the recipe uses the default parameters or those specified by the user when selecting the option "Use the parameters above as initial values in subsequent executions of this recipe".

10. **I have a data reduction problem different from the one listed in this tutorial. What I can do?** We have listed here only a limited number of data reduction problems we are aware of. We kindly ask the user to read the error message(s) and recipe log provided by reflex to try to understand the problem. If the origin and solution of the problem not clear and not already described in this tutorial, follow the next steps. You may run into a different problem, in which case:

    - Provide a short description of the problem.
    - Provide the error log if available (e.g. from reflex) or at least the full error message that is reported by the reflex widget reporting t hat an error occurred.
    - Indicate the recipe affected.
    - Include the full data set list (all ARCFILEs involved in the reduction of your DataSet).
    - If the problem is in the reflex GUI please provide a snapshot.
    - If the problem concerns data reduction quality or is an unexpected results.
    - Please collect all problems by category: installation, robustness, reduction quality, reflex, unexpected behaviour/result, improvement.
    - Please report data reduction problems to usd-help@eso.org.

10 Bibliography

Allart et al., 2022, Automatic model-based telluric correction for the ESPRESSO data reduction software. Model description and application to radial velocity computation, 2022 A&A, 666A, 196A. [6] 41

Forchì V., 2012, Reflex User Manual, VLT-MAN-ESO-19000-5037, Issue 0.7, *ftp://ftp.eso.org/pub/dfs/reflex/ReflexUserManual-3.1.pdf* 8, 44, 45

NIRPS Pipeline User Manual, NIRPS-4500-GEN-UM-211 ????, Issue 16.0

NIRPS: an adaptive-optics assisted radial velocity spectrograph to chase exoplanets around M-stars, Proceedings of the SPIE, Volume 10400, id. 1040018 15, pp. (2017). [7].[8]

7

[6] https://ui.adsabs.harvard.edu/abs/2022A%26A...666A.196A/abstract
[7] https://ui.adsabs.harvard.edu/abs/2017SPIE10400E..18W/abstract
[8] https://www.eso.org/public/teles-instr/lasilla/36/nirps/

A Useful information

A.1 Reflex directory naming conventions

We provide here a short description of the usual directory naming conventions, referring to the NIRPS case.

A standard `install_esoreflex` based installation, for example under `install_dir` will generate the following directory structure:

- `data_wkf`, this is the directory to which points the Reflex global variable `ROOT_DATA_DIR`. It contains the following directories:
  - `reflex_input`, containing the demo data, in the sub-directory espdr/raw. The global variable `RAW_DATA_DIR` points to `ROOT_DATA_DIR/reclex_input/espdr`
  - `reflex_book_keeping`, which will be populated by all the reflex bookeeping information, in the subdirectory espdr/. In particular the file bookkeeping.db, that is a mini-sql based database created, updated and used by reflex, the DataOrganizer/ directory, containing information on the current input data organisation classification and association, a set of directories called as `recipe_name_1` (for example `espdr_mflat_1`) which contain a set of subdirectories named as YYYY-MM-DDThh:mm:ss.mmm with the current execution timestamp id (for example 2019-04-16T21:42:54.335) and a link 'latest' pointing to the latest execution. These directories contains information relative to the esorex execution, and may be useful to debug if necessary. There are also other directories called (in case of NIRPS) as `InteractiveWorkflow_recipe_name`, `InteractiveWorkflow_espdr_mflat`
  - `reflex_tmp_products`, this directory contains under the subdirectory espdr/ the recipe temporary products (products of calibration recipe executed during the reduction chain) in subdirectories named like `recipe_name_1` and the directories `InteractiveWorkflow_recipe_name_1` (usually empty, but which may contain products of the python scripts used to handle the Reflex interactivity)
  - `reflex_logs` with a structure similar to `reflex_tmp_products` and containing the log of each recipe execution.
  - `reflex_end_products`, that contains a directory nammed with the current timestamp, as YYYY-MM-DDThh:mm:ss.mmm, containing a README file and other directories named as YYYY-MM-DDThh:mm:ss.mmm, by the timestamps of the first raw frame of the product they refer to, and containing all the final products of the corresponding execution.

A.2 Available Organisation Classification and Association (OCA) rules

The NIRPS workflow contains two OCA rules, the `nirps_wkf_full_chain.oca`, which carries out the execution of the full reduction chain (provided one has all required input data and static calibrations).

A.3 Glossary

We define here all relevant acronyms.

A.3.1 NIRPS data reduction naming conventions

We provide here the definition for the most frequent acronyms used in the tutorial. The user is also invited to read the pipeline user manual.

- S1D, extracted and merged spectrum 1D, flux (ADU or physical units) versus wavelength.
- S2D, combined detectors (BLUE and RED camera) extracted spectra order-by order, a 2D image containing at each raw the extracted spectrum of a given order. At the bottom of the image are stored orders corresponding to the BLUE camera, at the top orders corresponding to the red camera.
- CCF: Cross Correlation Function.
- SINGLEHR, single High Resolution observing mode. It may have two binning (XxY): 1x1 or 2x1.
- SINGLEUHR, single Ultra High Resolution observing mode. Only 1x1.
- MULTIMR, Multi-UT, Medium Resolution observing mode. It may have two binnings: 4x2, 8x4. It combines the light from any of the UTs.
- THAR: Thorium Argon wavelength calibration lamp.
- FP: Fabry Perot wavelength calibration unit.
- LFC: Laser Fiber Comb wavelength calibration unit.
- RO: (detector) Read Out region/port.
- RON: Read Out Noise.
- EXTi: Extension (index number i)
- FS2D: Order by order extracted spectrum
- BLAZE: Blaze spectrum for a given fibre
- `ORDER_PROFILE`: Order profile for a given fibre
- SNR: Signal to Noise Ratio
- CONAD: Conversion Analog to Digital factor [e^-/ADU]
- FSPECTRUM:Order by order extracted spectrum for a given fibre
- DLL: Distance-Wavelength.
- HOT: Hot Pixel (pixel with value significantly above the detector noise)
- BAD: Bad (not-linear) Pixel
- BERV: Barycentric Earth Radial Velocity.
- RV: Radial velocity

- BJD: Barycentric Julian date.
- DRIFT: relative shift of a wavelength calibration solution between the time of (night) observation and the time when the calibration was acquired (during the day).

A.3.2 Reflex workflow acronyms

- OCA: Organisation Classification and Association
- DataSet: data set
- SoF: set of frames