


AIDA



AIDA²⁰²⁰



EUDAQ User Manual

EUDAQ Development Team

Last update on October 2016
for EUDAQ version v1.7

This document provides an overview of the EUDAQ software framework, the data acquisition framework used also by the EUDET-type beam telescopes [1]. It describes how to install and run the DAQ system and use many of the included utility programs, and how users may integrate their DAQ systems into the EUDAQ framework by writing their own Producer – for integrating the data stream into the acquisition – and DataConverterPlugin – for converting data for offline analysis (e.g. the EUTelescope analysis framework).

Contents

1. License	4
2. Introduction	5
2.1. Architecture	5
2.2. Directory and File Structure	6
3. Installing EUDAQ	8
3.1. Installation of prerequisites	8
3.2. Download the source code from GitHub	11
3.3. Configuration via CMake	11
3.4. Compilation	13
4. Running EUDAQ	15
4.1. Preparation	15
4.2. Processes	18
4.3. Running the DAQ	24
4.4. Other Utilities	26
5. Writing a Producer	37
5.1. Configuration	37
5.2. Receiving Commands	37
5.3. Sending Data and the RawDataEvent class	39
5.4. Log Messages	40
5.5. Interfacing Python-Code via the PyProducer Interface	41
6. Data Conversion	42
6.1. StandardEvent and StandardPlane	42
6.2. LCIO and LCEvent	46
6.3. DataConverterPlugin	47
7. Other Parts of the Framework	49
7.1. FileWriter	49
7.2. FileReader	49
7.3. PluginManager	50
7.4. OptionParser	51
7.5. Timer	54
7.6. Utils	54
8. Reporting Issues	56
9. Developing and Contributing to EUDAQ	57
9.1. Regression Testing	57
9.2. Committing Code to the Main Repository	57
A. Source Code	59
A.1. Example Config File	59

A.2. Example Producer	60
A.3. Example DataConverterPlugin	67
A.4. Example Reader	70
B. Introduction to the build system and project files on Windows	72
B.1. MSBUILD	72
B.2. Project Files	72
B.3. Known Problems	75
C. Online Monitor Configuration Settings	76
C.1. Configuration Sections Overview	76
C.2. Configuration options in [General]	76
C.3. Configuration options in [Correlations]	76
C.4. Configuration options in [Clusterizer]	76
C.5. Configuration options in [HotPixelFinder]	76
C.6. Configuration options in [Mimosa26]	76
C.7. Configuration Example	77
Glossary	78

1. License

This program is free software: you can redistribute it and/or modify it under the terms of the Lesser GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the Lesser GNU General Public License for more details.

You should have received a copy of the Lesser GNU General Public License along with this program. If not, see <http://www.gnu.org/licenses/>.

2. Introduction

The EUDAQ software is a data acquisition framework, written in C++, and designed to be modular and portable, running on Linux, Mac OS X, and Windows. It was written primarily to run the EUDET-type beam telescope [2, 3], but is designed to be generally useful for other systems.

The hardware-specific parts are kept separate from the core, so that the core library can still be used independently. For example, hardware-specific parts are two components for the EUDET-type beam telescope: the Trigger Logic Unit (TLU) and the National Instrument system (NI) for Mimosa 26 sensor read out.

The raw data files generated by the DAQ can be converted to the Linear Collider I/O (LCIO) format, allowing for analysing the data using the EU Telescope package [4].

2.1. Architecture

It is split into a number of different processes, each communicating using TCP/IP sockets (compare Figure 1). A central Run Control provides an interface for controlling the whole DAQ system; other processes connect to the Run Control to receive commands and to report their status.

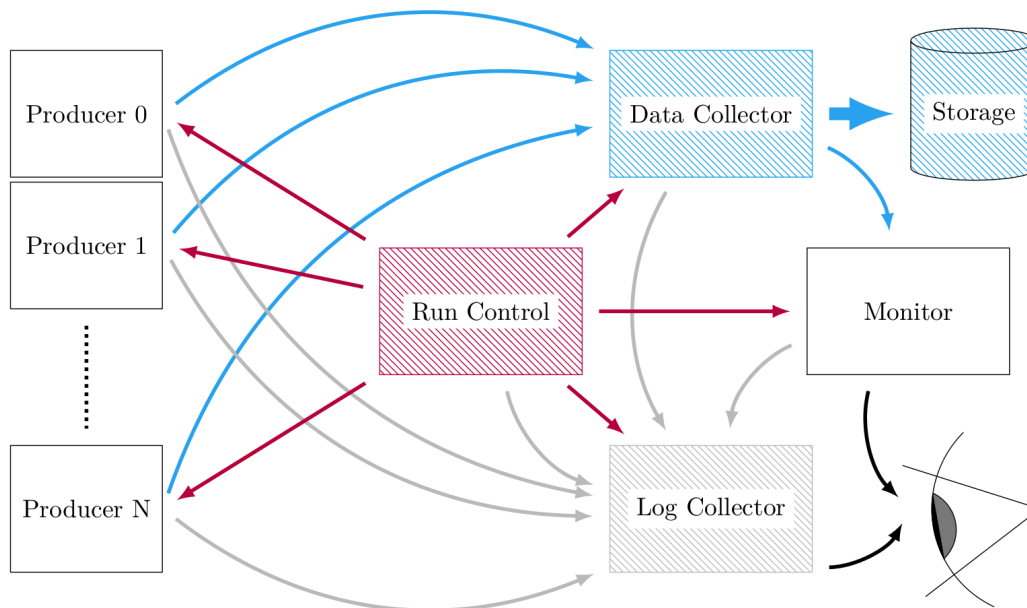


Figure 1: Schematic of the EUDAQ architecture [5].

Each hardware that produces data (e.g. the TLU, the NI, or a device under test (DUT)) will have a Producer process (on the left in Figure 1). This will initialize, configure, stop and start the hardware by receiving the commands from the Run Control (red arrows), read out the data and send it to the Data Collector (blue arrows).

The Data Collector receives all the data streams from all the Producers, and combines them into a single stream that is written to disk (Storage). It writes the data in a native raw binary format, but it can be configured to write in other formats, such as LCIO. The Log Collector receives log messages from all other processes (grey arrows), and displays them to the user, as well as writing them all to file. This allows for easier debugging, since all log messages are stored together in a central location. The Monitor reads the data file and generates online-monitoring plots for display. In the schematic it is shown to communicate with the Data Collector via a socket, but it actually just reads the data file from disk.

2.2. Directory and File Structure

The EUDAQ software is split into several parts that can each be compiled independently, and are kept in separate subdirectories. The general structure is outlined below:

- **main** contains the core EUDAQ library with the parts that are common to most of the software, and several command-line programs that depend only on this library. All definitions in the library should be inside the `eudaq` namespace. It is organised into the following subdirectories:
 - `lib/src` contains the library source code,
 - `exe/src` contains the (command line) executables source code,
 - `include` contains the header files inside the `eudaq` subdirectory (to match the namespace),
- **gui** contains the graphical programs that are built with Qt, such as the RunControl and LogCollector.
- **producers** contains all (user-provided) producers shipped with the EUDAQ distribution, for example:
 - `tlu` contain the parts that depend on the TLU.
 - `ni` contain the parts that depend on the NI system for Mimoso 26 read out.
 - e.g. `depfet`, `fortis`, `taki...` contain the code for third-party producers that have been used with EUDET-type beam telescopes.
- **extern** stores external software that is not part of EUDAQ itself, but that is needed by EUDAQ in some cases, such as the `ZestSC1` driver and the `tlufirmware` for the TLU.
- **bin** and **lib** contain the compiled binaries (executables and libraries) generated from the other directories.
- **conf** contains configuration files for running the beam telescope.
- **data** and **logs** are directories for storing the data and log files generated while running the DAQ.
- **doc** contains documentation, such as this manual.

Each directory containing code has its own `src` and `include` subdirectories, as well as a local `CMakeLists.txt` file containing the rules for building that directory using `CMake`. Header files usually have a `.hh` extension so that they can be automatically recognised as C++ (as opposed to C), and source files have either `.cc` for parts of a library or `.cxx` for executables.

Each directory can contain a `README.md` file for brief documentation for this specific part, e.g. as installation advice. Using the `*.md` file ending allows for applying the Markdown language [6]. Accordingly, content will be formatted on the the GitHub platform, where the code is hosted online.

3. Installing EUDAQ

The installation is described in four steps:¹

1. Installation of (required) prerequisites
2. Downloading the source code (GitHub)
3. Configuration of the code (CMake)
4. Compilation of the code

If you occur problems during the installation process, please have a look into the issue tracker on GitHub.² Here you can search, if your problem had already been experienced by someone else, or you can open a new issue (see section 8).

3.1. Installation of prerequisites

EUDAQ has few dependencies on other software, but some features do rely on other packages:

- To get the code and stay updated with the central repository on GitHub git is used.
- To configure the EUDAQ build process, the CMake cross-platform, open-source build system is used.
- To compile EUDAQ from source code requires a compiler that implements the C++11 standard.
- The libusb library is needed to communicate over USB with a TLU [7].
- Qt is required to build GUIs of the e.g. Run Control Log Collector.
- ROOT is required for the Online Monitor.

3.1.1. Git

Git is a free and open source distributed version control and is available for all of the usual platforms [8]. It allows for local version control and repositories, but also communicating with central online repositories like GitHub.

In order to get the EUDAQ code and stay updated with the central repository on GitHub git is used (see subsection 3.2). But also for developing the EUDAQ code having different versions (tags) or branches (development repositories), git is used (see section 9).

¹Quick installation instructions are also described on <http://eudaq.github.io/> or in the main README.md file of each branch, e.g. <https://github.com/eudaq/eudaq/blob/v1.6-dev/README.md>.

²Go to <https://github.com/eudaq/eudaq/issues>

3.1.2. CMake (required)

In order to generate configuration files for building EUDAQ (makefiles) independently from the compiler and the operating platform, the CMake build system is used.

CMake is available for all major operating systems from <http://www.cmake.org/cmake/resources/software.html>. On most Linux distributions, it can usually be installed via the built-in package manager (aptitude/apt-get/yum etc.) and on OSX using packages provided by e.g. the MacPorts or Fink projects.

3.1.3. C++11 compliant compiler (required)

The compilation of the EUDAQ source code requires a C++11 compliant compiler and has been tested with GCC (at least version 4.8), Clang (at least version 3.1), and MSVC (Visual Studio 2012 and later) on Linux, OS X and Windows.

If you are using Scientific Linux, please install the *Developer Toolset* available e.g. from <http://linux.web.cern.ch/linux/devtoolset/> to get access to a GCC version which fully implements C++11, e.g. on SL6 do

```
scl use devtoolset-1.1 bash
```

and cmake and install in this bash.

3.1.4. libusb (for the EUDET TLU)

In order to communicate over USB with a TLU, the libusb library is needed. Therefore, if you want to compile the `tlu` subdirectory, you should make sure that libusb is properly installed.

On Mac OS X, this can be installed using Fink or MacPorts. If using MacPorts you may also need to install the `libusb-compat` package.

On Linux it may already be installed, otherwise you should use the built-in package manager to install it. Make sure to get the development version, which may be named `libusb-devel` instead of simply `libusb`, e.g. on Ubuntu:

```
sudo apt-get install libusb-dev
```

On Windows, libusb is only needed if compiling with cygwin, in which case you should use the cygwin installer to install libusb. Otherwise libusb is not needed, as the included ZestSC1 libraries should work as they are.

3.1.5. ZestSC1 drivers and TLU firmware files (for the EUDET TLU)

Additionally to the libusb library, the EUDET TLU producer requires the ZestSC1 driver package and the FPGA firmware bitfiles. These are available to download via AFS from DESY. If AFS is accessible on the machine when CMake is running, the necessary files will be installed automatically. Otherwise, manually copy full folder with sub-directories from

- `/afs/desy.de/group/telescopes/tlu/ZestSC1` and
- `/afs/desy.de/group/telescopes/tlu/tlufirmware`

into the `extern` subfolder in your EUDAQ source directory.

3.1.6. Qt (for GUIs)

The graphical interface of EUDAQ uses the Qt graphical framework. In order to compile the `gui` subdirectory, you must therefore have Qt installed. It is available in most Linux distributions as the package `qt4-devel` or `qt5-devel` but make sure the version is at least 4.4, since there are a few issues with earlier versions.

If the included version is too old, or on other platforms, it can be downloaded from <http://qt.nokia.com/downloads>. Select the LGPL (free) version, then choose the complete development environment (it may also work with just the framework, but this is untested). Make sure the `QTDIR` environment variable is set to the Qt installation directory, and the `$QTDIR/bin` directory is in your path.

If you are using OSX, the easiest way to install Qt is using the packages provided by the MacPorts project (<http://www.macports.org/>).

3.1.7. ROOT (for Monitor)

The Online Monitor, as well as a few command-line utilities (contained in the `root` subdirectory), use the ROOT package for histogramming. It can be downloaded from <http://root.cern.ch> or installed via your favorite package manager.

Make sure ROOT's `bin` subdirectory is in your path, so that the `root-config` utility can be run. This can be done by sourcing the `thisroot.sh` (or `thisroot.ch` for csh-like shells) script in the `bin` directory of the ROOT installation:

```
source /path-to/root/bin/thisroot.sh
```

3.1.8. LCIO / EUTelescope (for converting/analysis)

To enable the writing of LCIO files, or the conversion of native files to LCIO format, EUDAQ must be linked against the LCIO and EUTelescope libraries. Detailed instructions on how to install both using the `ilcinstall` scripts can be found at <http://eutelescope.web.cern.ch/content/installation>.

The `EUTELESCOPE` and `LCIO` environment variables should be set to the installation directories of EUTelescope and LCIO respectively. This can be done by sourcing the `build_env.sh` script as follows:

```
source /path-to/Eutelescope/build_env.sh
```

3.2. Download the source code from GitHub

The EUDAQ source code is hosted on GitHub [9]. Here, we describe how to get the code and install a stable version release. In order to get information about the work flow of developing the EUDAQ code, please find the relevant information in see section 9.

3.2.1. Downloading the code (clone)

We recommend to obtain the software by using git, since this will allow you to easily update to newer versions. The source code can be downloaded with the following command:

```
git clone https://github.com/eudaq/eudaq.git eudaq
```

This will create the directory `eudaq`, and download the latest version into it.

Note: Alternatively and without version control, you can also download a zip/tar.gz file of EUDAQ releases (tags) from <https://github.com/eudaq/eudaq/releases>. By downloading the code, you can skip the next two subsections.

3.2.2. Changing to a release version (checkout)

After cloning the code from GitHub, your local EUDAQ version is on the master branch (check with `git status`). For using EUDAQ without development or for production environments (e.g. at test beams), we strongly recommend to use the latest release version. Use

```
git tag
```

in the repository to find the newest stable version as the last entry. In order to change to this version in your local repository, execute e.g.

```
git checkout v1.6.0
```

to change to version v1.6.0.

3.2.3. Updating the code (fetch)

If you want to update your local code, e.g to get the newest release versions, execute in the `eudaq` directory:

```
git fetch
```

and check for new versions with `git tag`.

3.3. Configuration via CMake

CMake supports out-of-source configurations and generates building files for compilation (makefiles). Enter the `build` directory and run CMake, i.e.

```
cd build
cmake ..
```

CMake automatically searches for required packages and verifies that all dependencies are met using the `CMakeLists.txt` scripts in the main folder and in all sub directories. By default, only the central shared library, the main executables and (if Qt4 or Qt5 have been found) the graphical user interface (GUI) are configured for compilation. You can modify this default behavior by passing the `-DBUILD_[name]` option to CMake where `[name]` refers to an optional component, e.g.

```
cmake -DBUILD_gui=OFF -DBUILD_tlu=ON ..
```

to disable the GUI but enable additionally executables of the TLU producer. Find some of the most important building options in Table 1.

option	default	comment
<code>-DBUILD_main</code>	ON	Builds main EUDAQ executables. The common library, and some command-line programs that depend on only this library
<code>-DBUILD_manual</code>	OFF	Builds Manual in pdf-format. In Ubuntu e.g.: <code>pdflatex</code> , <code>scrartcl.cls</code> and <code>upquote.sty</code> are required, thus execute <code>sudo apt-get install texlive-latex-base texlive-latex-recommended texlive-latex-extra</code>
<code>-DBUILD_tlu</code>	OFF	Builds the TLU producer and command-line tools. <code>libus</code> , <code>ZestSC1</code> and <code>tlufirmware</code> files are required.
<code>-DBUILD_<producername></code>	OFF	<code>-DBUILD_<producername>=ON</code> is needed to enable building of a specific producer. The <code>producername</code> is the same as the name of the producer's directory in <code>./producer</code> . These are user-contributed producers for specific detectors operating with the EUDET-type beam telescope. They should not be compiled unless needed.
<code>-DBUILD_gui</code>	ON	Builds GUI executables, such as the Run Control and Log Collector. Requires QT4/5.
<code>-DBUILD_python</code>	ON	Builds Python EUDAQ binding library.
<code>-DBUILD_pybindgen</code>	OFF	Builds pybindgen binding libraries.
<code>-DBUILD_onlinemon</code>	ON	Builds Online Monitor executable. Requires ROOT.
<code>-DBUILD_offlinemon</code>	OFF	Builds offline monitor executable. Requires ROOT.
<code>-DBUILD_metamon</code>	OFF	Builds MetaData Monitor executable.
<code>-DBUILD_resender</code>	OFF	Builds resender producer.
<code>-DBUILD_nreader</code>	OFF	Builds native reader Marlin processor used for data conversion into LCIO. Requires LCIO/EUTelescope
<code>-DINSTALL_PREFIX=<PATH></code>	<code><eudaq></code>	In order to install the executables into <code>bin</code> and the library into <code>lib</code> of a specific <code><path></code> , instead of into the <code><eudaq></code> path. The corresponding de-installation can be done by: <code>cd <eudaq>build && sudo xargs rm <install_manifest.txt</code>

Table 1: Some of the most important building options for CMake.

Note: After generating building files by running `cmake ..`, you can list all possible option and their status by running `cmake -L`. Using a GUI version of CMake shows also all of

the possible options.

Corresponding settings are cached, thus they will be used again next time CMake is running. If you encounter a problem during installation, it is recommended to clean the cache by just removing all files from the build folder, since it only contains automatically generated files. In order to start from scratch, just run:

```
cd build
rm -rf *
```

3.4. Compilation

3.4.1. Compilation on Linux/OSX

From the top EUDAQ directory, run the command

```
cd build
make install
```

in order to compile the common library with some command-line programs (the contents of the `./main/exe` subdirectory). If other parts are needed, you can specify them as arguments to the CMake command during the configuration step (see 3.3).

The executable binaries and the common shared library will be installed by default into the `bin` and `lib` directories in the source tree, respectively. If you would like to install into a different location, please set the respective parameter during the CMake configuration.

3.4.2. Setup and Compilation on Windows using Visual Studio

This section gives a short overview on the steps needed to compile the project under Windows (tested under Windows 7, 32-bit and 64-bit). For a more detailed introduction to the Windows build system and Visual Studio project files see the appendix B on page 72.

- Prerequisites
 - Compiler: Download Visual Studio Express Desktop 2012 or later (e.g. 2013 Version): <http://www.microsoft.com/en-us/download/details.aspx?id=40787>
 - Download Qt4 or later Qt5
 - Download and install the pthreads library (pre-build binary from <ftp://sources.redhat.com/pub/pthreads-win32>) into either `c:\pthreads-w32` or `./extern/pthreads-w32`
- Start the Visual Studio *Developer Command Prompt* from the Start Menu entries for Visual Studio (Tools subfolder) which opens a `cmd.exe` session with the necessary environment variables already set. If your Qt installation has not been added to the

global %PATH% variable, you need to execute the `qtenv2.bat` batch file (or similar) in the Qt folder, e.g.

```
C:\Qt\Qt5.1.1\5.1.1\msvc2012\bin\qtenv2.bat
```

Replace "5.1.1" with the version string of your Qt installation.

- Now clone the EUDAQ repository (or download using GitHub) and enter the build directory on the prompt, e.g. by entering

```
cd c:\Users\[username]\Documents\GitHub\eudaq\build
```

- Configuration: Now enter

```
cmake ..
```

to generate the VS project files.

- Compile by calling

```
MSBUILD.exe EUDAQ.sln /p:Configuration=Release
```

or install into `eudaq\bin` by running

```
MSBUILD.exe INSTALL.vcxproj /p:Configuration=Release
```

- This will compile the main library and the GUI. For additional processors, please check the individual documentation.

Note on "*moc.exe - System Error: The program can't start because MSVCP110.dll is missing from your computer*" errors: when using Visual Express 2013 and `pthread-w32` 2.9.1, you might require "Visual C++ Redistributable for Visual Studio 2012": download (either x86 or x64) from <http://www.microsoft.com/en-us/download/details.aspx?id=30679> and install.

4. Running EUDAQ

This section will describe running the DAQ system, mainly from the point of view of EUDET-type beam telescope [10] operated together with DUTs. However, this description can be applied to DAQ system in general.

All executable programs from the different subdirectories are placed inside the `bin` subdirectory, and should be run from here. They should all accept a `-h` (or `--help`) command-line parameter, which will provide a summary of possible different command-line options.

The executable programs mainly split up in two different categories: Processes, which are used for the data acquisition and communicating with the Run Control (DAQ), and utilities, which are used before or after the data taking in order to access the data files (Test, Devel., Tools). In Table 2, you will find an overview of the most important EUDAQ executables.

4.1. Preparation

Some preparation is needed to make sure the environment is set up correctly and the necessary TCP ports are not blocked before the DAQ can run properly.

4.1.1. Directories

EUDAQ expects two directories that will be used to store data files and log files. These can be directories or symbolic links to other directories.

Firstly, inside the `eudaq` directory, there should be a directory (or a symbolic link) called `data`. This will contain the data files written by the Data Collector, as well as a file containing the last run number, so that it will continue incrementing even when the DAQ is restarted. Secondly, there should be a directory (or symbolic link) called `logs`. This will be used by the Log Collector to store log files containing all the log messages received.

4.1.2. Hostnames

EUDAQ processes communicate between themselves using TCP/IP sockets. When processes are started, they need to know where the Run Control runs. There is no completely fool-proof way of determining this, so processes look at the environment variable `$HOSTNAME`.

Usually, this should be the DNS name of the machine it is running on, but in some cases it may not work correctly. If this is the case, it may be necessary to set this variable manually, either to the real host name or the machine's IP address. In the case, that all the processes will be running on the same computer, the host name can be set to `localhost` which corresponds to the IP address `127.0.0.1`.

Depending on the operating system and the command shell, you can set the host name by

Category	Name	Binary	Description
DAQ	Run Control	euRun.exe	GUI version, recommended (Sec. 4.2.1)
DAQ	Run Control	TestRunControl.exe	CLI version (Sec. 4.2.1)
DAQ	Log Collector	euLog.exe	GUI version, recommended (Sec. 4.2.2)
DAQ	Log Collector	TestLogCollector.exe	CLI version (Sec. 4.2.2)
DAQ	Data Collector	TestDataCollector.exe	CLI, recommended (Sec. 4.2.3)
DAQ	Online Monitor	OnlineMon.exe	GUI version, recommended (Sec. 4.2.6)
DAQ	Online Monitor	TestMonitor.exe	CLI version (Sec. 4.2.6)
DAQ	TLU Producer	TLUProducer.exe	CLI for EUDET-type telescopes (Sec. 4.2.4)
DAQ	Mimosa26 Producer	NIProducer.exe	CLI for EUDET-type telescopes (Sec. 4.2.5)
DAQ	Run Listener	RunListener.exe	CLI, listens to Run Control, no description
Test	TLU Control	TLUControl.exe	CLI (Sec. 4.4.1)
Test/DAQ	Test Producer	TestProducer.exe	CLI (Sec. 4.2.7)
Test/DAQ	Dummy Producer	euProd.exe	GUI, no description
Devel./DAQ	Example Producer	ExampleProducer.exe	CLI, recommended (Sec. 4.2.7)
Devel.	Exampler Reader	ExampleReader.exe	CLI, example raw file reader, no description
Devel.	Option Example	OptionExample.exe	CLI, example to use options, no description
Tools	File Checker	FileChecker.exe	CLI, raw data file checker (Sec. 4.4.2)
Tools	Test Reader	TestReader.exe	CLI, raw data file reader (Sec. 4.4.3)
Tools	EURunSplitter	EURunSplitter.exe	CLI, splits raw data files, no description
Tools	Converter	Converter.exe	CLI, raw data file converter (Sec. 4.4.4)
Tools	IPHC Converter	IPHCConverter.exe	CLI, IPHC M26 to raw converter, no description
Tools	Cluster Extractor	ClusterExtractor.exe	CLI, extracting clusters from raw files (Sec. 4.4.5)
Tools	MagicLogBook	MagicLogBook.exe	obsolete (!), CLI (Sec. 4.4.6)

Table 2: Overview of EUDAQ executables: DAQ processes and tools as graphical user interfaces (GUI) or command line interfaces (CLI).

- for bash-like shells: `export HOSTNAME=name`
- for csh-like shells: `setenv HOSTNAME name`
- for Windows command lines / scripts: `set HOSTNAME=name`

where `name` is the name or the IP address.

Note: It is recommended to set the host name to the (local) IP address. This method is approved and working at the EUDET-type telescopes at DESY and CERN.

4.1.3. Ports and firewall

The different processes communicate between themselves using TCP/IP sockets. If a firewall is running, it may block these connections, especially if the processes are running on different computers. If all the processes will be run from the same computer, then it is probably not necessary to do anything. If a port is blocked, you will see an error message similar to the following when attempting to start some programs:

```
Are you sure the server is running? - Error 61 connecting to ↔
localhost:44000: Connection refused
```

The ports can be configured when calling the the processors on the command line (see below), but the default and usually free port numbers are:

44000: This is the port used to send commands from the Run Control.

44001: This port is used to send data from the producers to the Data Collector.

44002: This port is used to send log messages from all processes to the Log Collector.

If processes will be running on different computers, then these ports should be opened up in the firewall. The method for doing this depends on the Operating System used, and is outside the scope of this manual.

4.1.4. TLU permissions

If you are not using a TLU connected to a Linux OS, you may skip this part.

On many Linux distributions, the device node used to communicate over the USB bus is only accessible by a user having root rights by default (`sudo ...`). To set the correct permissions when a TLU is connected, you need to add a `udev` rule: as a root user, create the file `/etc/udev/rules.d/54-tlu.rules` and add the following lines:

```
# for Debian
ACTION=="add", DRIVERS=="?*", ATTR{idVendor}=="165d", ↔
ATTR{idProduct}=="0001", MODE="0666"
```

in case you are using a debian-based distribution such as Ubuntu.

```
# for Red Hat, e.g. SL5
SYSFS{idVendor}=="165d", SYSFS{idProduct}=="0001", GROUP="NOROOTUSB", ↵
MODE="0666"
```

if you are using a Red Hat-based distribution (such as Scientific Linux) or:
After replugging the TLU, the device should be accessible by all users.

4.2. Processes

The DAQ system is made up of a number of different processes that may all be run on the same, or on different computers.

4.2.1. Run Control

There are two versions of the Run Control – a text-based version and a graphical version (see Figure 2). The graphical version is recommended, since it is well tested and complete. The executable is called `euRun.exe`, or on Mac OS X it is an application bundle called `euRun.app`. The text-based version can be useful for testing, the executable is `TestRunControl.exe`.




Figure 2: The Run Control graphical user interface.

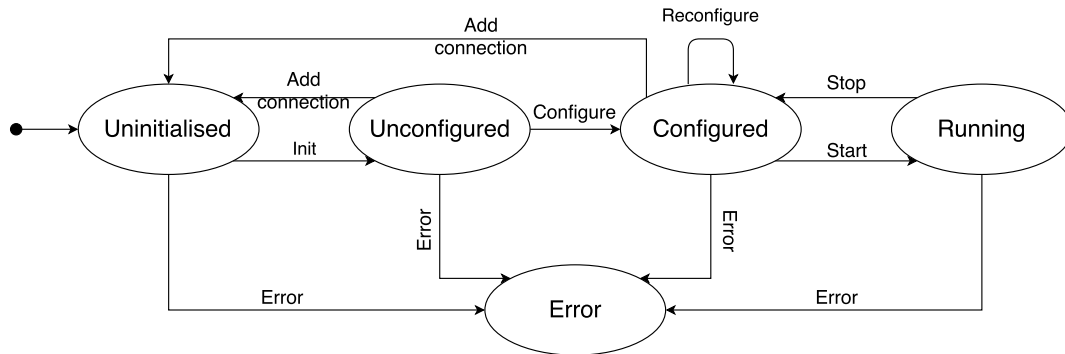


Figure 3: The FSM of EUDAQ.

Usually, no command-line option should be needed. If needed, it can be told to listen on a specific port (e.g. to run two copies on the same machine) using the `-a <port>` option, for example:

```
./euRun.exe -a 44000
```

Note: If two copies of EUDAQ should run simultaneously, the ports of the Log and DataCollectors have to be different!

Finite-State Machine Since EUDAQ version 1.7, a finite-state machine (FSM) is implemented (see Figure 3) [11]. Each process connected to the Run Control can always be characterized by the current state:

- **UNINITIALISED**: the initial state of every connection. Initialisation has not been conducted yet. Available transitions are `OnInitialise` and `OnTerminate`.³
- **UNCONFIGURED**: initialisation has already been conducted, but configuration parameters have not been set yet. Available transitions are `OnConfigure` and `OnTerminate`.
- **CONFIGURED**: configuration parameters have already been set. Available transitions are `OnStartRun`, `OnConfigure` and `OnTerminate`.
- **RUNNING**: the state of the operating connections. For producers it means producing data, the `OnlineMonitor` are providing plots, the `LogCollector` is continuing

³Some detectors require not only setting up different parameters, but also the first initialisation of the hardware. Previously both these steps were carried out during the configuration. However, setting up the hardware is only necessary during the start and can be time demanding, so reinitialisation of the hardware can take additional time when reconfiguration is needed. One should note that `OnInitialise` function that is used to set the connection into the CONFIGURED state has default implementation in the base `CommandReceiver` class, so users who do not need an initialisation step can simply skip it. In addition, this ensures backward compatibility for existing producers. [11]

to collect log messages and the `DataCollector` combines data streams from producers. The only available transition is `OnStopRun`.

- `ERROR`: this state can be used by users in case of errors during configuration, running process, etc. The only available transition is `OnTerminate`.

The state of the machine is determined by the lowest state of the connected components (`LogCollector`, `DataCollector`, `OnlineMonotor`, `Producers`) in the following priority: `ERROR`, `UNITIALISED`, `UNCONFIGURED`, `CONFIGURED`, `RUNNING`. It means, for example, that even if only one connection is in the `ERROR` state, the whole machine will also be in that state. This prevents such mistakes as running the system before every component has finished the configuration.

4.2.2. Log Collector

It is recommended to start the Log Collector directly after having started the Run Control and before starting other processors in order to collect all log messages generated by all other processes.

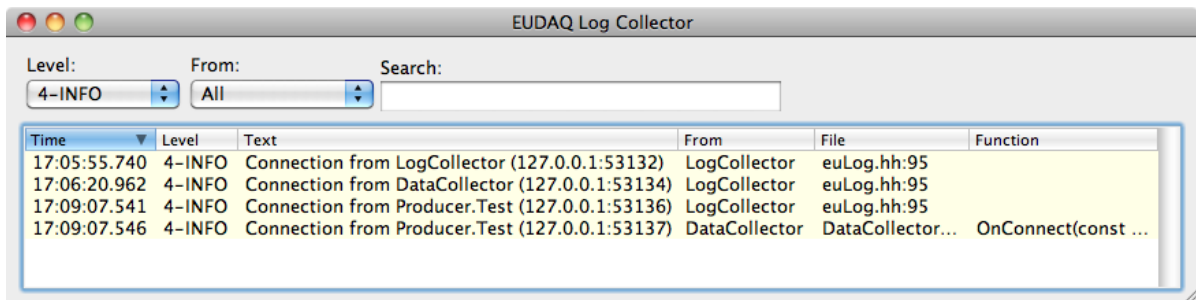


Figure 4: The Log Collector graphical user interface.

Like the Run Control, there are also two versions of the Log Collector. The graphical version is called `euLog.exe`, or `euLog.app` on Mac OS X, and the text-based version is called `TestLogCollector.exe`.

If it is running on the same machine as the Run Control, it should not need any command-line options. However, if it is running on a different machine, it must be told on which machine the Run Control is running. By using the `-r <hostname>` option, the Log Collector knows where to connect to the Run Control, for example the Run Control runs on a machine having a local IP 192.168.0.1 and listening on the port 40000 (as default):

```
./euLog.exe -r tcp://192.168.0.1:44000
```

The port can also be set using the `-a <port>` option, similar to the Run Control.

4.2.3. Data Collector

The Data Collector is the process that collects all the raw data from the Producers, merges all the connected incoming streams into a single data stream, and writes it to file. There is only a text-based version called `TestDataCollector.exe`. It is recommended to start the Data Collector directly after having started the Run Control and RunControl and before starting other processors.

Like the Log Collector, it should be told where to connect to the Run Control if it is not running on the same machine. Accordingly, the `-r` and `-a` options can be used, for example:

```
./TestDataCollector.exe -r tcp://192.168.0.1:44000
```

It is also possible to run multiple Data Collector instances within one EUDAQ session. This can be useful to reduce network traffic and e.g. write the output of one producer to a locally attached disk. When running several Data Collectors simultaneously, the Run Control assigns a Producer to a Data Collector by name: if the name of a Data Collector matches that of a Producer, the latter will be given the address and port of the former. There can be only one instance of an *unnamed* Data Collector which serves as the default for any non-matching Producer; if no unnamed Data Collector is present, the first one connecting will serve as the default.

The name of a Data Collector can be set with the `-n` option, for example:

```
./TestDataCollector.exe -n myproducer
```

If you wish to run several instances of the Data Collector on one machine, you need to make sure that they listen to different addresses using the `-a` option as described above. Furthermore, each Data Collector has to write to a different file by including the `FilePattern` option in the corresponding section of your configuration file (also see section 4.3.3):

```
[DataCollector.myproducer]
FilePattern = "../data/run$6R_myproducer$X"
```

4.2.4. TLU Producer

If you do not have a TLU in your setup, you may skip this part. Otherwise you should run a TLU Producer, which will configure the TLU, and read out the timestamps and send them to the Data Collector. On the computer with the TLU connected, start the `TLUProducer.exe` program. If this is not the same machine as the Run Control, use the `-r` option as for the Data and Log Collectors, in this example:

```
./TLUProducer.exe -r 192.168.0.1:44000
```

If the TLU Producer fails to start, make sure the permissions are set up correctly (see subsection 4.1.4).

4.2.5. NI Producer (Mimosa 26 sensors)

If you don't have a EUDET-type beam telescope using 6 planes with Mimosa 26 sensors [3] in your setup, you may skip this part. The NI Producer is the interface between the Mimosa DAQ (at the moment Anemome 1.3 LabView software) and the EUDAQ. The Mimosa DAQ is based on a National Instrument machine which runs with a Windows OS. Thus, you start the NI Producer only on a NI crate after opening the Anemome LabView software. It will connect to the Run Control (euRun) by using `-r` option, too, in this example:

```
./NIProducer.exe -r 192.168.0.1:44000
```

If the NI Producer fails to start, make sure that the IP addresses (host names) are set correctly in the start scripts and the EUDAQ configurations files [10, section 5.2.1.2 1].

4.2.6. Online Monitor

The Online Monitor reads the data file written by the Data Collector, and generates several ROOT histograms that can be useful for online monitoring. Since it reads the native data file directly (by using the corresponding DataConverterPlugin), it must be run on the same machine as the Data Collector.

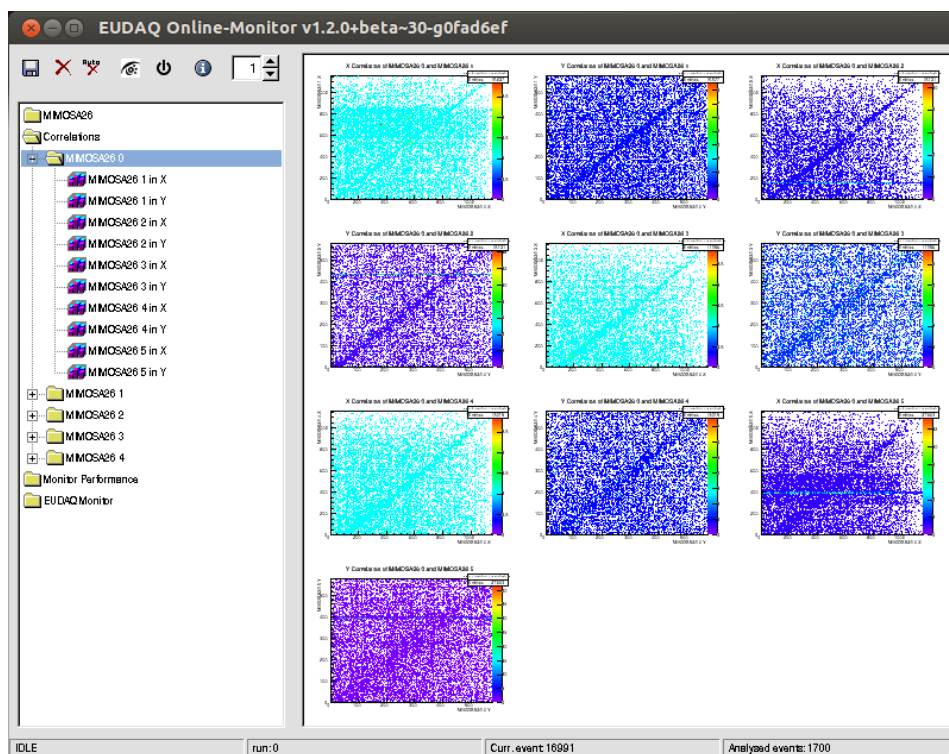


Figure 5: The OnlineMon showing correlation plots between different Mimosa26 planes of the EUDET telescope.

The Online Monitor can be run in one of two modes: online or offline. In online mode, it connects to the Run Control, so it will know when new runs are started, and it will automatically open each new data file as it is created. To run it in online mode, the `-r` option may be used to assign Run Control, in this example:

```
./OnlineMon.exe -r 192.168.0.1:44000
```

Note: The Online Monitor is working properly on Unix machines. In addition, it is recommended to run the Online Monitor on the same machine as the Data Collector. In offline mode, there is no Run Control, and it only analyses the data file it is given on the command line using the `-f` option. An example command line is:

```
./OnlineMon.exe -f 5432
```

This will run it in offline mode, opening the file corresponding to run 5432 (alternatively, the full path to a file may be given).

4.2.7. Test and Example Producer

For testing purposes, you may use the TestProducer. This works similarly to a real producer, but does not talk to any real hardware, instead providing a menu for the user to manually send events.

The ExampleProducer was written to illustrate the writing of a new Producer (see section 5). However, it will actually generate some example data, and so can also be used for testing purposes. It works more like a real Producer than the TestProducer, in that it does not require user intervention to generate each trigger, and the data generated emulates a simple (but realistic) sensor, and can be properly converted, and therefore displayed in the Monitor.

4.2.8. Other/DUT Producer(s)

If you have a producer for your own hardware (see section 5), it should also have an option to set the address of the Run Control.

4.2.9. Slow Producer

Since EUDAQ version 1.7 [11]: A slow producer is a special type of producers which interacts with the `DataCollector` in a different way using the hardware conception of the triggerless data taking. Instead of sending busy signals it can simply provide data at its own rate. The `DataCollector` distinguishes the slow producer from the usual one, it waits only for events from simple producers and ignores the absence of those from slow producers. In order to integrate a slow device one has to simply extend it from the `SlowProducer` class instead of the `Producer` class, but having then the same functionality.

4.2.10. Python Interface and Wrapper for Core EUDAQ Components

A Python interface is provided for selected EUDAQ components: RunControl, DataCollector and a Producer, that can be extended on the Python side. The interface is realized through the `ctypes` package that is part of every standard Python installation and requires the `numpy` Python package to be installed. The interface code for all components is located in the `main/python` directory.

To use the interface and access the components as Python objects, the wrapper must be loaded inside your Python script:

```
#!/usr/bin/env python2
execfile('PyEUDAQWrapper.py') # load ctypes wrapper

prc = PyRunControl() # start run control with default settings
# wait for more than one active connection to appear
while prc.NumConnections < 2:
    sleep(1)
prc.Configure("ExampleConfig") # load configuration file
while not prc.AllOk:
    sleep(1) # sleep while waiting for all connected producers
prc.StartRun()
```

This little script creates a RunControl instance, sends a configuration to all connected producers, waits for their reply, and starts a new run. Several more extensive examples for using Python with EUDAQ are located in the `python` directory in the main EUDAQ directory.

4.3. Running the DAQ

4.3.1. Starting EUDAQ using STARTRUN scripts

To start EUDAQ, all of the necessary processes have to be started in the correct order. The first process must be the Run Control (`euRun`), since all other processes will attempt to connect to it when they start up. Then it is recommended to start the Log Collector, since any log messages it receives may be useful to help with debugging in case everything does not start as expected. Next, the Data Collector should be started. Finally all other Producers, and if needed, the Online Monitor.

In the `eudaq/etc/scripts` directory, there are different Unix `STARTRUN` scripts. These scripts can be customized to load the appropriate processes for running the DAQ. This allows you to start all the processes necessary with a single command.⁴

Starting scripts for Unix and for Windows, you can find in the `eudaq-configuration` repository [12].

⁴If starting processes on other computers via SSH, it is recommended to set up SSH keys so that the processes may be started without having to type a password.

4.3.2. Operating EUDAQ

Once all the processes have been started, and by using the Run Control (see Figure 2), According to the machine state section 4.2.1, EUDAQ and all processes can be initialized, configured or re-configured, data taking (runs) can be started and stopped, and the software can be terminated.

- First, the appropriate initialisation file should be selected (see subsection 4.3.3 for creating and editing init-files). Then the **Init** button can be pressed, which will send a initialisation command to all connected processes.
- Second, the appropriate configuration should be selected (see subsection 4.3.3 for creating and editing configurations), and the **GeoID** should be verified (see subsection 4.3.4), before continuing. Then the **Config** button can be pressed, which will send a configuration command (with the contents of the selected configuration file) to all connected processes. The full contents of the configuration file will also be stored in the beginning-of-run-event (BORE) of the data file, so that this information is always available along with the data.
- Once all connected processes are fully configured, a run may be started, by pressing the **Start** button. Whatever text is in the corresponding text box ("Run:") when the button is pressed will be stored as a comment in the data file. This can be used to help identify the different runs later.
- Once a run is completed, it may be stopped by pressing the **Stop** button. Runs will also stop and restart automatically when the data file reaches a threshold in size (by default this is 1 GB).⁵ The threshold size for restarting a run may be configured in the config file (see subsection 4.3.3).
- At any time, a message may be sent to the log file by filling in the ("Log:") text box and pressing the corresponding button. The text should appear in the LogCollector window, and will be stored in the log file for later access.
- Once the run is stopped, the system may be reconfigured with a different configuration, or another run may be started; or EUDAQ can be terminated.

4.3.3. Init/Config-Files

*.init-files for initialisation and *.conf-files for configuration are text files in a specific format, containing name-value pairs separated into different sections. See subsection A.1 for an example file.

Any text from a # character until the end of the line is treated as a comment, and ignored. Each section in the config file is delimited by a name in square brackets (e.g. [RunControl]). The name represents the type of process to which it applies; if there

⁵This is because there is a file size limit of 2 GB for storage on the GRID, and the processed files can grow bigger than the original native files.

are several such processes, then they can be differentiated by including the name after a period (e.g. [Producer.Example]). Within each section, any number of parameters may be specified, in the form `Name = Value`. It is then up to the individual processes how these parameters are interpreted.

The entire contents of the config file will be sent to all processes during the configuration, and each process will have the appropriate section selected. The file will also be attached to the BORE, so that it is available with the data later, even if the original config file is modified or deleted.

4.3.4. GeoID

The GeoID is a number representing the physical positioning of the telescope and DUT(s). Each time a change is made to the telescope layout, this number should be incremented. To change the number, double-click on it, and a window will appear with the new value. By default it will increment the old value by one, so normally you should just click OK, but if necessary you may edit the value first.

The GeoID is inserted into the config file when it is sent, so it is also stored in the data file, and will be used to select the correct GEAR file for alignment during the data analysis stage.

4.4. Other Utilities

There are a number of other utilities available that are not needed for running the DAQ, but can be useful for other tasks such as debugging. The executables are all located in the `bin` subdirectory. They should all accept a help (`-h` or `--help`) option, to print a summary of the available options.

4.4.1. TLUControl

The `TLUControl.exe` program is a standalone program for running the TLU without using the full DAQ. The most commonly used parameters are shown below. For each option, the short (preceded by one dash) and the long (preceded by two dashes) option names are shown. Only one of the two forms should be used for each option, but long and short options can be mixed together on the command line) Each options has a parameter and a default value that will be used if the option is not specified.

- `-d --dutmask <mask = 0>`: The DUT mask; this defines which DUT connections are activated. It is a bit-mask, so 1 means connector 0, 2 means connector 1, etc..
- `-a --andmask <mask = 255>`: The AND mask; this defines which external trigger inputs are activated. It is a bit-mask, so 1 means channel 0, 2 means channel 1, etc.. The specified channels are ANDed together, and used to generate a trigger signal.
- `-t --trigger <msecs = 0>`: Internal trigger period. If non-zero, the TLU will generate internal triggers with the specified period in milliseconds. If set to zero, the internal trigger is off.

- i `--dutinputs` `<values = ">`: Input mode select. A sequence of comma-separated strings specifying which connectors to use for the DUT inputs. Valid values are RJ45, LEMO, HDMI, and NONE.
- u `--wait-for-user`: Pause the program after the TLU is configured, before starting triggers. The default is to not wait for the user.

Other parameters available are as follows:

- o `--ormask` `<mask = 0>`: The OR mask; this defines which external trigger inputs are activated. It is a bit-mask, so 1 means channel 0, 2 means channel 1, etc.. The specified channels are ORed together, and used to generate a trigger signal.
- v `--vetomask` `<mask = 0>`: The VETO mask; this defines which external trigger inputs are activated. It is a bit-mask, so 1 means channel 0, 2 means channel 1, etc.. The specified channels are used to veto the generation of a trigger if they are active.
- w `--wait` `<ms = 1000>`: Wait time. This is the time to wait between updates.
- n `--notimestamp`: Indicates that the timestamp buffer should not be read out.
- q `--quit`: Quit the program after configuring the TLU.
- s `--save-file` `<filename = ">`: The filename to save trigger numbers and timestamps
- p `--strobeperiod` `<cycles = 1000>`: Period for timing strobe (in TLU clock cycles).
- l `--strobelength` `<cycles = 100>`: Length of 'on' time for timing strobe (in TLU clock cycles).
- b `--dutveto` `<mask = 0>`: Mask for enabling veto of triggers ('backpressure') by raising DUT_CLK.
- hm `--handshakemode` `<nohandshake = 0>`: In this mode the TLU issues a fixed-length pulse on the trigger line (0 = no handshake).
- pw `--powervctrl` `<mV = 800>`: [obsolete but provided for backward compatibility, please use -pv] Sets the Vcntl control voltage to all PMTs. The range of values is between 0 and 1000 (or 0 and 2000 if the TLU has been modified by cutting LC1 and jumpering L01 on the PMT Supply Daughterboard and specifying the -pm 1 option).
- pv `--pmtvcntl` `<mV = 800>`: Sets the Vcntl control voltage to all PMTs (see option -pw for more details). Will override the value of -pw if it is specified. If neither -pw or -pv is specified, the default value will be used (and can be overridden on an individual PMT basis).

- p1 --pmtvcntl1 <mV>: Sets the PMT Vcntl voltage for PMT1 (Chan 0) only. If not specified, the default or values specified by -pw or -pv (which will override -pw) is used.
- p2 --pmtvcntl2 <mV>: Sets the PMT Vcntl voltage for PMT2 (Chan 1) only. If not specified, the default or values specified by -pw or -pv (which will override -pw) is used.
- p3 --pmtvcntl3 <mV>: Sets the PMT Vcntl voltage for PMT3 (Chan 2) only. If not specified, the default or values specified by -pw or -pv (which will override -pw) is used.
- p4 --pmtvcntl4 <mV>: Sets the PMT Vcntl voltage for PMT4 (Chan 3) only. If not specified, the default or values specified by -pw or -pv (which will override -pw) is used.
- pm --pmtvcntlmod <value = 0>: Specifies whether the TLU PMT Supply Daughter-card is modified (LC1 cut and L01 jumpered) or not. A <value> of 0 specifies that it is unmodified (and thus the Vcntl range is from 0mV to 1000mV), and a <value> of 1 specifies that the TLU is modified (and thus the Vcntl range is from 0mV to 2000mV). This feature is to accomodate newer Hamamatsu PMT models (e.g. H10721) that require a control voltage range of, for instance, 500mV to 1100mV that are being used in place of the older (discontinued, but what the TLU was designed to accomodate and control) models that required a control voltage of between 250mV and 900mV.
- f --bitfile <filename = "">: The bitfile containing the TLU firmware to be loaded.
- e --error-handler <value = 2>: Error handler setting. Setting to 0 indicates the program should abort on an error. Setting it to a value greater than 0 indicates the number of tries that should be attempted before generating an exception.
- r --fwversion <value = 0>: Specifies the firmware version to load (setting to 0 indicates the version should be chosen automatically).
- z --trace-file <filename = "">: The filename to save a trace of all USB accesses. Prepend a dash ('-') to output errors only, or a plus ('+') for all data (including block transfers).

An example use of the command is shown below:

```
./TLUControl.exe -t 200 -d 3 -i LEMO,RJ45 -u
Using options:
TLU version = 0 (auto)
Bit file name = '' (auto)
Trigger interval = 200 ms (5 Hz)
DUT Mask = 0x03 (3)
```

```
Veto Mask = 0x00 (0)
And Mask  = 0xff (255)
Or Mask   = 0x00 (0)
DUT inputs = LEMO,RJ45
Strobe period = 0x0003e8 (1000)
Strobe length = 0x000064 (100)
Enable DUT Veto = 0x00 (0)
Save file = '' (none)
```

```
TLU Version = v0.2c
TLU Serial number = 0x062b (1579)
Firmware file = TLU2_Toplevel.bit
Firmware version = 65
Library version = 65
```

Press enter to start triggers.

TLU Started!

```
Status: 20,00,--,--,--,-- (0,0)
Scalers: 0, 0, 0, 0
Particles: 2
Triggers: 0
Entries: 0
TS errors: 0, 0 (redundancy, re-read)
Timestamp: 0x8d768 (579432) = 0.00150891
Time: 0.009 s, Freq: 0 Hz, Average: 0 Hz
```

```
0, 0x27fb479 (41923705) = 0.109174, diff=41923705
1, 0x7139ab9 (118725305) = 0.309174, diff=76801600
2, 0xba780f9 (195526905) = 0.509174, diff=76801600
3, 0x103b6739 (272328505) = 0.709174, diff=76801600
4, 0x14cf4d79 (349130105) = 0.909174, diff=76801600
```

```
Status: 20,00,--,--,--,-- (0,1)
Scalers: 0, 0, 0, 0
Particles: 7
Triggers: 5
Entries: 5
TS errors: 0, 0 (redundancy, re-read)
Timestamp: 0x1726fa48 (388430408) = 1.01152
Time: 1.023 s, Freq: 4.92913 Hz, Average: 4.88442 Hz
```

```
5, 0x196333b9 (425931705) = 1.10917, diff=76801600
6, 0x1df719f9 (502733305) = 1.30917, diff=76801600
7, 0x228b0039 (579534905) = 1.50917, diff=76801600
8, 0x271ee679 (656336505) = 1.70917, diff=76801600
```

```

          9, 0x2bb2ccb9 (733138105) = 1.90917, diff=76801600
Status:   20,00,--,--,--,-- (0,1)
Scalers:  0, 0, 0, 0
Particles: 12
Triggers: 10
Entries:   5
TS errors: 0, 0 (redundancy, re-read)
Timestamp: 0x2e5bb708 (777762568) = 2.02538
Time: 2.037 s, Freq: 4.93259 Hz, Average: 4.90838 Hz
^CQuitting...

```

This sets up internal triggers at 5 Hz (200 ms period), and activates DUT inputs 0 and 1. Input 0 is configured to use the LEMO connector, and input 1 to use the RJ45 connector. The first part of the output just summarizes the input parameters. The next part shows information about the version numbers of the TLU and the firmware.

It will then configure the TLU, and if the `-u` option is used, it will wait for the user to press enter before continuing. The triggers are then enabled, and a summary of the status is printed out periodically (by default every 1 second). The program can be stopped cleanly by pressing `Ctrl-C`.

Each block of status output consists of:

- a list of triggers, if there were any since the last update (the first time there are none), each showing:
 - the trigger number,
 - the timestamp of the trigger, in hex, decimal and converted to seconds,
 - the difference since the last trigger.
- the status of the DUT connections (see below),
- the values of the scalers on the external trigger inputs,
- the number of “particles”, which means all the potential triggers (including those that were vetoed),
- the number of triggers that actually got sent to the DUTs,
- the number of entries in the trigger buffer, this should be equal to the number of triggers printed out at the top of the status block,
- the number of timestamp errors detected by redundancy, and by re-reading,
- the current timestamp value,
- the time since the run started, the current trigger frequency, and the average frequency over the whole run.

In the example output this block is repeated three times, before `Ctrl-C` is pressed to stop it. The status is of the DUT connections formatted as:

- two digits for each DUT connection consisting of:
 - two hyphens (--) if the connection is inactive, else

- the first digit represents the inputs from the DUT; with the busy line in bit 0 and the clock line in bit 1 (note the clock input can float low or high if a LEMO input is selected, as it is not connected),
- the second digit represents the state of the FSM, as defined in the TLU manual[7] (0 is ready, 1 is waiting for busy high, 4 is waiting for busy low, 5 is DUT-initiated veto, and F is an error condition).
- then in parentheses:
 - the veto state (software veto in bit 0, overall veto in bit 1),
 - the DMA state (1 when a DMA transfer is taking place).

4.4.2. FileChecker

This is a small utility that reads raw data files and checks if all events are readable, can be synchronised using the TLU trigger id and lists which type of subevents the file contains.

It should be called with list of file paths or run numbers. For any argument that consist only of numerical digits the file path is constructed by substituting `$6R` in the input pattern (defaults to “`./data/run$6R.raw`”) with the run number padded to 6 digits. For example:

```
./FileChecker.exe {6045..6050}
```

This would produce the following output.

run	valid	num_events	contains	errors
6045	true	13131	MUPIX4,NI,TLU	
6046	true	1	MUPIX4,NI,TLU	
6047	true	14674	MUPIX4,NI,TLU	
6048	true	7776	MUPIX4,NI,TLU	
6049	false	0		no events in the file.
6050	false	-1		read error.

4.4.3. TestReader

The `TestReader.exe` program will read a native data file, and can display various pieces of information from the file. Commonly used options are:

- `-b`: Display the BORE.
- `-e`: Display the end-of-run-event (EORE).
- `-d <range>`: Display the specified range of event numbers.
- `-p`: Process the displayed events and display the corresponding StandardEvents.

- u**: Dump the raw data for the displayed events.
- s**: Try to resynchronize events based on the TLU event number. A full description of this option is outside the scope of this manual (but if you don't know what it is, you probably don't need it).

After the options a list of one or more filenames can be given. Any filenames that consist only of numerical digits will be interpreted according to the input pattern (by default this is “`../data/run$6R.raw`”, where `$6R` will be replaced with the run number padded to 6 digits). For example:

```
./TestReader.exe -b -e -p -d 1-10,100,1000 example.raw 5432
```

This will display the BORE and EORE, and the events 1 to 10, 100 and 1000, processing them to also display the StandardEvents, from the files `example.raw` and `../data/run005432.raw`.

4.4.4. Converter

The `Converter.exe` program will read a native data file, optionally select just a subset of events from the file, and can then write it out to another file in either the same native format, or a different format. The most commonly used options are:

- t** `<type>`: The file type to write out. The available types are listed below.
- e** `<range>`: Select the specified range of event numbers.
- s**: Try to resynchronize events based on the TLU event number (see `TestReader` in subsection 4.4.3).

The available output file types are as follows:

native: The native EUDAQ binary file format, consisting of a serialised stream of `DetectorEvents`, containing the raw data read out from the hardware.

standard: Like the `native` format, this is also a serialised stream, but in this case it contains `StandardEvents`, in which the raw data has been converted into a standard format.

lcio: The standard LCIO file format used by the analysis software. This type is only available if EUDAQ was compiled with LCIO support.

root: A Root file containing a TTree with the hit pixel information.

text: A simple text based format (not yet implemented).

mimoloop: A text based format mimicking the output of the `mimoloop` program (from Angelo Cotta Ramusino and Lorenzo Chiarelli at INFN Ferrara).

Although this program can be used to convert a native data file into LCIO format, the more usual (and therefore better tested) way is to use the `EUTelescope` converter.

4.4.5. ClusterExtractor

This program can be used to quickly extract some clusters from raw data. It is not as sophisticated as the EUTelescope package, which should be preferred for real analysis, but it can be useful for doing quick checks. It will read a native data file, perform a basic clustering, and then write these clusters to one text file per sensor plane. The most commonly used options are:

- p `<pixels>`: The cluster size in pixels. It should be an odd number, with 1 meaning no clustering (just pixels over threshold), 3 meaning 3×3 pixel clusters, etc.
- n `<adcs>`: The noise level (sigma) in ADC units. This is used to scale the thresholds in terms of the noise.
- s `<thresh>`: The threshold for seed pixels, in terms of the noise.
- c `<thresh>`: The threshold for the total charge of a cluster, in terms of the cumulative noise of all the pixels in the cluster.
- w: Reports the cluster centre as the weighted average of the pixels, instead of the position of the seed pixel.

An example use is:

```
./ClusterExtractor.exe -p 3 -n 3.5 -s 6 -c 10 -w 5432
```

This will generate a number of text files named `runNNN_eutel_M.txt`, where `NNN` is the run number, and `M` is the sensor plane number. The format of the output text files is as follows:

```

2      2      51487659237
 182   153    126
 241   120    125
3      1      51489095892
 111   67     346
5      1      51491334074
 113   141    171
7      2      51495330212
 252   240    305
 95    170    189
```

The first line contains the event number, the number of clusters, and the TLU timestamp. Then for each cluster there is one line, containing the `x` and `y` coordinates of the cluster centre, and the total charge in ADC units. The cluster lines are prepended with a space to make it easier to scan the file by eye.

4.4.6. MagicLogBook

This program is designed to extract as much information as possible from data files and log files, in order to reconstruct a log book. Despite its name, it is in fact not magical, so it is preferable to keep a good log book during running, rather than relying on this program to generate it later.

The available options are listed below:

- f **<fields>**: A list of fields to include in the output, in the form **name=value**, with multiple fields separated by commas. If a predefined list is also specified these will be appended to the list.
- s **<separator>**: The separator to use between fields in the output. The default is a tab character.
- h **<string>**: A string that appears at the beginning of the header line (with the list of field names), that can be used to differentiate it from the other lines. The default is an empty string.
- p **name**: Use a predefined list of fields. Currently available values are **normal** and **full**.
- o **<file>**: The output filename. By default the standard output is used.

The easiest method of running is to use a predefined list of fields. There are currently two predefined lists available: **normal** and **full**. If neither of these are suitable, contact the EUDAQ maintainer, as it may be possible to add more options.

The **normal** list includes:

- the run number,
- the config file name,
- the run start time,
- for the teawts (EUDRBs):
 - the mode,
 - the sensor type,
 - whether they are running unsynchronized,
 - the number of boards,
 - and the firmware version.
- and for the TLU:
 - the internal trigger interval,
 - the AND mask,
 - the DUT mask,
 - and the firmware version.

The `full` list includes all the values from the `normal` list, plus the number of events in the run and the end of run time. This is because these values can only be known by reading the whole data file to the end, which is slow, especially for large data files.

If necessary, other information is available using custom fields, although the syntax for these is a bit complicated, since it is designed to be as flexible as possible at specifying any information in the data file. In the future it may be redefined in order to simplify it if possible. Therefore it is recommended to use a predefined list of fields where possible. Custom fields are specified as a comma separated list of items in the form `name=value`, with the name being what will appear on the header line of the output, and the value specifying what exactly to extract from the file. The possible values are illustrated below, although not exhaustively:

`events*`: The number of events in the run.

`config`: The configuration name, or:

`config:section:key`: The value of the `key` from the corresponding `section` in the config (e.g. `config:Producer.EUDRB:NumBoards`).

`bore`, `tlu`, `eudrb`, `eore*`: Something from the BORE, the TLUEvent or EUDRBEvent subevents of the BORE, or the EORE, respectively:

`bore:.Run`: The run number

`bore:(name)`: Otherwise, if the second part does not start with a period, the value of the tag `(name)` is used (e.g. `tlu:DutMask` or `eudrb:MODE`).

`log`: Something from the log file (not implemented yet).

* items marked with an asterisk require reading the whole data file, and are therefore slow, especially when large data files are involved.

Note that the `EUDRBEvent` is now deprecated, having been replaced by the `RawDataEvent`, but there is currently no way to specify this.

The `MagicLogBook` command is used as follows:

```
./MagicLogBook.exe -p normal ../data/*.raw
```

This will produce an output similar to the following:

Run	Config	Mode	Det	Start	U	P	Trg	AND	DUT	Tfw	Efw
6371	eudet-beam			2009-07-29 07:44:39.535	1	6	0	0xf	0x10	241	
6372	eudet-beam			2009-07-29 08:03:05.079	1	6	0	0xf	0x10	241	
6373	eudet-m26test			2009-07-30 09:57:45.157	1	6	255	0xff	0x12	241	
6374	eudet-m26test			2009-07-30 10:00:45.205	1	6	255	0xff	0x12	241	
6375	eudet-m26test			2009-07-30 10:05:38.625	1	6	1	0xff	0x12	241	
6376	eudet-m26test			2009-07-30 10:10:00.107	1	6	1	0xff	0x12	241	
6379	eudet-m26test			2009-07-30 10:13:05.322	1	6	1	0xff	0x12	241	

Note that the header row has been modified slightly to fit into the page width: the `U` should be `UnSync`, `P` should be `Planes`, `Trg` should be `TriggerInterval`, `Tfw` should be `TLUfw`, and `Efw` should be `EUDRBfw`. The columns `Mode`, `Det` and `EUDRBfw` are missing

from the output due to the fact that this information is now stored in a `RawDataEvent`, which is not currently accessible with this version of the program.

4.4.7. Others

Some programs that are less used (or recently added) may not be described here. If they look interesting, you can find out more about them by running them with the help (`-h` or `--help`) option, or by examining the source code.

5. Writing a Producer

Producers are the binding part between a user DAQ and the central EUDAQ Run Control. A Producer consists of a `CommandReceiver` and a `DataSender`, where the first receives commands from the Run Control while the latter allows to send binary data events to the Data Collector. A Producer base class is provided in order to simplify the integration. The Producer is compiled against the EUDAQ library which implements all of the required communication over the network. Example code for producers is provided, see subsection A.2.

5.1. Configuration

The `Configuration` class is a way of storing configuration information in a way that is easily accessible, and can be saved to or loaded from a human-readable file (see subsection 4.3.3), and can be sent over the network. It is defined in the following header:

```
#include "eudaq/Configuration.hh"
```

The configuration consists of a number of sections, each of which contains a list of name-value pairs. The values are stored as strings, but they can be converted to/from arbitrary types. Methods are provided to load from or save to file, to set the current section, and to set or get configuration values. An example use is shown below:

```
std::ifstream infile("../conf/ExampleConfig.conf");
eudaq::Configuration config(infile, "Producer.Example");
int param = config.Get("Parameter", 0);
std::cout << "Loaded config, param = " << param << std::endl;
config.Set("Parameter", param+1);
config.Set("OtherParam", "something");
std::ofstream outfile("Test.conf");
config.Save(outfile);
```

This creates a configuration loaded from the file `../conf/ExampleConfig.conf`, selecting the `Producer.Example` section. It then gets an integer parameter from the configuration and displays it. Then it modifies the value of the parameter and sets another parameter, before writing the configuration to the file `Test.conf`.

A configuration object will be received by the Producer during the configuration, as described in subsection 5.2.2.

5.2. Receiving Commands

The Run Control distributes commands to all registered clients, controlling the global finite state machine and thus the DAQ. Whenever a user input is received from the Run Control, the corresponding member function of the Producer will be invoked. These member functions of the Producer class are virtual and can be overloaded by the user

in their respective Producer implementations. The `Producer` base class definition is provided in the header file:

```
#include "eudaq/Producer.hh"
```

5.2.1. OnInitialise

This method is called whenever an initialise command is received from the Run Control. The method signature is:

```
virtual void OnInitialise(const eudaq::Configuration & config);
```

The configuration object is received from the Run Control. The Producer is free to read any part of the configuration object but parameters cannot be changed owing to the `const` qualifier.

If the initialisation step is not required for a certain DAQ, the function can be left unimplemented in the Producer. In that case the corresponding member function of the `CommandReceiver` base class is executed.

5.2.2. OnConfigure

This method is called whenever a configure command is received from the Run Control. The method signature is:

```
virtual void OnConfigure(const eudaq::Configuration & config);
```

The configuration object is received from the Run Control where it has been loaded from a user configurable file. The Producer is free to read any part of the configuration object but parameters cannot be changed owing to the `const` qualifier.

5.2.3. OnStartRun

This is called on the start of each run. The method signature is:

```
virtual void OnStartRun(unsigned param);
```

As a parameter, it receives the run number of the started run. The Producer must send a BORE, and then enable the data acquisition on the attached devices and prepare for receiving events from the hardware.

5.2.4. OnStopRun

This is called at the end of the run. The method signature is simply:

```
virtual void OnStopRun();
```

Care should be taken that there are no more events pending to be read out. Once all data events have been sent, an EORE should also be sent, to signal to the DAQ that the Producer has ended the run successfully.

5.3. Sending Data and the RawDataEvent class

Events may be sent to the DAQ using the Producer's `SendEvent()` method that has the following signature:

```
void SendEvent(const Event &);
```

It takes as a parameter an object derived from the `eudaq::Event` base class that will be serialised and sent to the Data Collector. In practice it will usually be of the type `RawDataEvent`.

The `RawDataEvent` is a generic container for blocks of binary data and is used to encapsulate the data read directly from the devices and send it to the central DAQ for storage. Each `RawDataEvent` may contain any number of raw data blocks. By convention each block usually corresponds to one sensor, but this is not required; it is up to each Producer how the raw data are encoded. The `RawDataEvent` class is defined in the following header file:

```
#include "eudaq/RawDataEvent.hh"
```

The class is described in the following in more detail.

5.3.1. Constructor

A `RawDataEvent` is constructed as follows:

```
RawDataEvent event("EXAMPLE", run, event);
```

Where "EXAMPLE" is a string unique to the particular producer that will be used by the event decoding factory to select the correct converter during decoding. The `run` and `event` parameters are the run number and event number, respectively.

Furthermore, the producer is required to send a BORE and EORE at the beginning and end of a run respectively. These are regular `RawDataEvent` objects with a particular flag set. For convenience, separate methods are available to create the respective event types:

```
RawDataEvent::BORE("EXAMPLE", run);
RawDataEvent::EORE("EXAMPLE", run, event);
```

These methods return a `RawDataEvent` that may be either be sent directly to the DAQ, or be modified first, e.g. by setting tags as described below in subsection 5.3.3.

5.3.2. Adding Data

Once a `RawDataEvent` has been constructed, data blocks may be added either using a vector:

```
std::vector<unsigned char> buffer = ...;
event.AddBlock(id, buffer);
```

or using a pointer to a block of memory, and a length in bytes:

```
unsigned char * buffer = ...;
event.AddBlock(id, buffer, len);
```

Where `id` is an integer used to uniquely identify the different blocks. The `buffer` variable points to the actual data for the block under consideration block and can be provided either as STL `vector`, where the full vector is read and sent, or as C-pointer to a memory block where the `len` argument is required to specify the size.

The type of the vector or data array can be freely chosen by the user since the data serializers are implemented as templates. However, if data types larger than `char` are used, special care has to be taken for appropriate endianness of the data. This is especially true if the producer and decoding software run on different machines which might have different architectures.

5.3.3. Tags

`RawDataEvent` objects as well as other types inheriting from the `Event` base class also provide the option to store tags. Tags are name-value pairs containing additional information which does not qualify as regular DAQ data which is written in the binary blocks. Particularly in the BORE this is very useful to store information about the exact sensor configuration which might be required in order to be able to decode the raw data stored. A tag is stored as follows:

```
event.SetTag("Temperature", 42);
```

The value corresponding to the tag can be set as an arbitrary type (in this case an integer), it will be converted to a STL string internally.

5.4. Log Messages

The logging infrastructure allows to send information, error messages or debug information to a central point in the DAQ system to collect logging information, the Log Collector. It is strongly encouraged to use the logging system rather than simple `cout` statements. The logging class is defined in:

```
#include "eudaq/Logger.hh"
```

The following macros for sending log messages are defined, listed here in decreasing order of severity:

EUDAQ_USER: A user-generated message (e.g. from the RunControl Log button).

EUDAQ_ERROR: Something that went wrong and should be fixed. Errors usually are blocking, i.e. the data acquisition cannot be continued without fixing the cause.

EUDAQ_WARN: A warning that something may not be quite right and should probably be taken care of. Warnings are considered non-blocking, i.e. the data acquisition

will proceed but some of the components might experience problems (lacking configurations for a threshold setting would be an example).

EUDAQ_INFO: An message generated during normal running containing information that may be useful to the user.

EUDAQ_EXTRA: Some extra information that may be less useful in normal running.

EUDAQ_DEBUG: Information for debugging purposes that will normally be hidden and should only be used for development purposes. Additional information for shifters should be categorized as **EUDAQ_EXTRA**.

They are used as follows:

```
EUDAQ_ERROR("No keyboard detected: press F1 to continue.");
```

The messages will be sent to the central Log Collector if it is connected, otherwise they will be displayed on the local terminal. The log level can be changed in the following way:

```
EUDAQ_LOG_LEVEL("WARN");
```

Any messages lower than the specified level will just be ignored. This can be useful to filter out unimportant messages and, for example, just display error messages.

5.5. Interfacing Python-Code via the PyProducer Interface

As described in section 4.2.10, a Python interface is provided for selected EUDAQ components including a Producer. This basic implementation can be extended on the Python-side as demonstrated by the example `python/example-producer.py`.

6. Data Conversion

Data are stored on disk, by default, in a native binary format, containing the raw data as read out by the various Producers. It is the same format used for serialising the data over the socket connection to the Data Collector. To be legible to other software components, this data must be converted into a standardised format so that the monitoring and analysis software does not require specific information about the functionalities and data encoding scheme of every detector, but can be applied generically to any sensor.

Currently, two different formats are provided for this purpose. The first is the `StandardEvent` type, an EUDAQ-internal class used by the online monitoring tool and other utilities of the framework. It is a very simplified format tailored towards pixelated tracking detectors. The second type is the LCIO standard format from the linear collider community, which is also used by the EUTelescope analysis software.

6.1. StandardEvent and StandardPlane

The `StandardEvent` is a class designed to represent pixel sensor data in a reasonably easy to use way, but still be flexible enough to store the data from a wide range of different sensors.

Each `StandardEvent` represents one event of data from the whole telescope and any DUTs, so a run will consist of a sequence of `StandardEvents`. It inherits from the `Event` base class, meaning that it has a run number, an event number, an optional timestamp, and may also contain tags (see subsection 5.3.3).

The decoded pixel data is stored in an array of `StandardPlanes`, each representing one sensor plane of the telescope or DUT. Each `StandardPlane` contains the charge values from the pixels of one sensor, and may contain several frames in cases where the sensor is read out multiple times per event. It also has the concept of a “result” frame, which is calculated from the one or more of the source frames according to different rules that may be specified with flags. The result frame contains only one charge value per pixel, with a positive signal, and is what will be used for the analysis. It may consist of either differences between the original frames (e.g. in the case of correlated double sampling (CDS)), a sum of all original frames, or specific parts of the different frames selected according to the pivot information. Flags may be set to select which of the different methods is used. It may also contain a submatrix number per pixel, which can be used to differentiate different parts of the sensor, so that they may be analyzed separately later, and a pivot boolean (true or false) per pixel, which can be used to indicate whether the pixel was sampled before or after the trigger, and is used to determine which parts of the sensor to combine when the `FLAG_NEEDCDS` flag is set.

Both the `StandardEvent` and the `StandardPlane` classes are defined in the following header file:

```
#include "eudaq/StandardEvent.hh"
```

In general, a user should not need to construct a `StandardEvent` object, but should create one or more `StandardPlanes`, that will be added to a given `StandardEvent`.

6.1.1. Constructor

The `StandardPlane` constructor has the following signature:

```
StandardPlane(unsigned id, const std::string & type,
              const std::string & sensor = "");
```

Where `id` is an arbitrary numerical identifier for the plane that can be used to differentiate between different planes of the same type, `type` is the type of the Producer that generated the frame (should be the same as that in the `Producer` and the `DataConverterPlugin`), and `sensor` is the name of the sensor, in the case that the Producer can read out more than one type of sensor.

6.1.2. SetSizeRaw and SetSizeZS

Once a `StandardPlane` has been constructed, the size should be set. There are two methods for doing this, depending on whether the data are stored in raw or zero-suppressed mode. In raw mode all pixels are stored, whether they have a signal or not. In zero-suppressed mode, only those with a signal above a certain threshold are stored, along with their coordinates, and any below the threshold are suppressed.

The signature of the `SetSizeRaw` method is:

```
void SetSizeRaw(unsigned w, unsigned h, unsigned frames = 1, int flags = 0);
```

Where `w` is the full width of the sensor (in the x-direction, usually columns) in pixels, `h` is the full height of the sensor (in the y-direction, usually rows) in pixels, `frames` is the number of frames, and `flags` may be a combination of the following values, separated by a bitwise OR (i.e. `|`):

`FLAG_NEEDCDS`: Indicates that the data are in 2 or 3 frames and that neighbouring frames should be subtracted to produce the result.

`FLAG_NEGATIVE`: Indicates that the charge values are negative, so should be negated to produce the result.

`FLAG_ACCUMULATE`: Indicates that all frames should be summed to produce the result.

`FLAG_WITHPIVOT`: Indicates that pivot information is stored per pixel, and should be used for constructing the result.

`FLAG_WITHSUBMAT`: Indicates that submatrix information is stored per pixel.

`FLAG_DIFFCOORDS`: Indicates that each frame can have different coordinates, in the case of zero-suppressed data, otherwise all frames will share the same coordinates.

The signature of the `SetSizeZS` method is as follows:

```
void SetSizeZS(unsigned w, unsigned h, unsigned npix,
               unsigned frames = 1, int flags = 0);
```

Where all parameters are the same as in `SetSizeRaw`, but there is an extra parameter (`npix`) that specifies how many pixels to preallocate. If the number of pixels above threshold is known, this may be used to allocate them all at once. If not, then this parameter may be set to zero, and pixels can be allocated as needed (but note that this way may be slower, since memory will need to be reallocated for each new pixel).

6.1.3. SetPixel and PushPixel

Once the size has been set, the values of the pixels can then be loaded into the `StandardPlane`. There are two methods for doing this: `SetPixel`, that sets the value of an already allocated pixel, and `PushPixel` that allocates space for a new pixel and sets that.

The signatures of `SetPixel` are as follows:

```
void SetPixel(unsigned index, unsigned x, unsigned y, unsigned pix,
              bool pivot = false, unsigned frame = 0);
void SetPixel(unsigned index, unsigned x, unsigned y, unsigned pix,
              unsigned frame);
```

where `index` is the index of the pixel to set, `x` and `y` are the coordinates of the pixel, and `pix` is the charge value for the pixel. The value of the pivot, and the frame number may optionally be set also, if relevant. Note that if only the pivot is set, care should be taken that it is of type `bool` to avoid accidentally setting the frame instead.

The signatures of `PushPixel` are as follows:

```
void PushPixel(unsigned x, unsigned y, unsigned pix,
               bool pivot = false, unsigned frame = 0);
void PushPixel(unsigned x, unsigned y, unsigned pix,
               unsigned frame);
```

where all parameters are the same as in `SetPixel`. The only difference being the lack of an `index` parameter, since this will always allocate a new pixel and append it to the existing list.

6.1.4. Setting other information

Other than the pixel values, the `StandardPlane` also stores some other information that should be set if applicable:

```
void SetTLUEvent(unsigned ev);
```

This sets the trigger ID as read out from the TLU. If it was read out and stored, it should be set using this method to allow cross checks in the analysis.

```
void SetPivotPixel(unsigned p);
```

This sets the value of the pivot pixel (or pivot row etc. – the value is arbitrary). It is only here to allow cross-checks in the analysis; if the pixels are to be combined using

the pivot information, then it should also be set in the per-pixel pivot values. The value here cannot be used for that purpose since the order of reading out the pixels is not in general known.

```
void SetFlags(FLAGS flags);
```

Some flags may be set after calling `SetSizeRaw` or `SetSizeZS`, but this is not possible with the flags `FLAG_WITHPIVOT`, `FLAG_WITHSUBMAT` or `FLAG_DIFFCOORDS` since these flags affect how memory is allocated by those methods.

6.1.5. Adding Planes to the StandardEvent

Once the plane has been constructed and filled, it may be added to a `StandardPlane` using the following method:

```
StandardPlane & AddPlane(const StandardPlane &);
```

This will copy the plane into the list of `StandardPlanes` stored by the `StandardEvent`. It will return a reference to the copy of the plane, that can be used to make further modifications if necessary.

6.1.6. Extracting information

The `StandardEvent` inherits the following methods from the `Event` base class:

```
unsigned GetRunNumber() const;
unsigned GetEventNumber() const;
uint64_t GetTimestamp() const;
T GetTag(const std::string & name, T def) const;
```

allowing access to the run number, event number, timestamp (if set) and any tags (where `T` is an arbitrary type). It also has the following methods to access the `StandardPlanes` that it contains:

```
size_t NumPlanes() const;
const StandardPlane & GetPlane(size_t i) const;
```

These return the number of planes stored, and a reference to a particular plane, respectively. The individual planes can then be examined using the following methods:

```
const std::string & Type() const;
const std::string & Sensor() const;
unsigned ID() const;
unsigned TLUEvent() const;
unsigned PivotPixel() const;
```

These return the type of the plane (i.e. the type of `Producer` / `DataConverter` that generated it), the type of sensor for the plane (in the case that the plane type can hold different types of sensor data), the ID of the plane (used to differentiate different planes

of the same type), the TLU trigger ID for the plane (if it was read out and stored) and the value of the pivot pixel (or pivot row) for the plane. Further information about the plane is available in:

```
unsigned XSize() const;
unsigned YSize() const;
unsigned NumFrames() const;
unsigned TotalPixels() const;
unsigned HitPixels() const;
unsigned HitPixels(unsigned frame) const;
```

These return the full width and height of the sensor in pixels, the number of frames stored for the plane, total number of pixels for the plane (i.e. full width \times height), the number of pixels over threshold (for zero-suppressed data) in the result frame, and the number of pixels over threshold in a particular source frame.

Note that for the `HitPixels` method, there are two versions; the first takes no parameter and returns the number of hit pixels in the result frame, while the second takes the frame number as a parameter and returns the number of hit pixels in that frame from the underlying source data. Normally the first version would be used, unless access is needed to the raw data from the sensor. Similarly, the other methods for accessing the data all have two versions:

```
double GetPixel(unsigned index) const;
double GetX(unsigned index) const;
double GetY(unsigned index) const;
const std::vector<pixel_t> & PixVector() const;
const std::vector<coord_t> & XVector() const;
const std::vector<coord_t> & YVector() const;
```

These return the charge value, the x coordinate and the y coordinate of a particular pixel (for the first three methods), or a vector of these values for all pixels in the frame (for the final three methods).

Here, `coord_t` and `pixel_t` are both `double`, even though the values stored are usually integers. This is in order to make the `StandardPlane` as general as possible, allowing it to store, for example, clusters with non-integer coordinates instead of pixels, and it also makes it easier to pass the values directly into Root histograms without first having to convert them to `double`. All the above methods also have a version taking the frame number (as the second parameter if they already have one parameter), which returns the information from the underlying source frame instead of the result frame.

6.2. LCIO and LCEvent

Another option available with the framework is the output of data as LCIO events. The LCIO format is a very flexible container defined and used by the linear collider community. The exact encoding of the data relies upon the requirements from the user and cannot be described in a generic way.

Many tracking detectors rely on the additional classes provided by the EUTelescope data analysis framework, examples for implementations can be found in the respective converter plugins in the EUDAQ source tree.

6.3. DataConverterPlugin

In order to allow different DUTs to easily incorporate their data into the monitoring and analysis chain, the `DataConverterPlugin` system was developed. This allows all the conversion code for each producer to be kept in one file, with the necessary parts being called automatically as needed. This section describes how to write a new converter plugin, to use existing converter plugins see subsection 7.3.

Writing a converter plugin for a new producer involves defining a new class that derives from the `DataConverterPlugin` base class and implementing a few methods. Each converter plugin contains a unique string that defines which type of `RawDataEvents` it is able to convert. This is the same string that is set in the `RawDataEvent` when it is created by the relevant producer. The `DataConverterPlugin` class is defined in the following header:

```
#include "eudaq/DataConverterPlugin.hh"
```

The methods to be implemented are described below, and a full example is provided in subsection A.3.

6.3.1. Constructor

The constructor should call the `DataConverterPlugin` constructor, and pass as a parameter the string representing the type of `RawDataEvent` this plugin can convert. A single static instance of the converter should then be defined, and instantiated in the source file. This is illustrated below:

```
class ExampleConverterPlugin : public eudaq::DataConverterPlugin {
    ExampleConverterPlugin() : eudaq::DataConverterPlugin("EXAMPLE") {
        // constructor...
    }
    // more methods...
    static ExampleConverterPlugin m_instance;
};
ExampleConverterPlugin ExampleConverterPlugin::m_instance;
```

this will cause the constructor to be called during initialization of the program, and the `DataConverterPlugin` constructor will automatically register the plugin and make it available in the `PluginManager`.

6.3.2. Initialization

Every time a new run is started, the `Initialize` method will be called. It has the following signature:

```
virtual void Initialize(const Event & ev, const Configuration & c);
```

It receives as parameters the BORE, and the configuration used for the run. The plugin may extract any tags from the BORE, or other information from the configuration, and store it in member variables for use during decoding.

6.3.3. GetTriggerID

Since each producer that reads out the trigger ID from the TLU stores it differently in the raw data, there is no general way to extract this information. The `GetTriggerID` method remedies this, by providing a generic interface to access the trigger ID. The signature is as follows:

```
virtual unsigned GetTriggerID(const Event & ev) const;
```

It receives the `Event` as a parameter, from which it should extract the TLU trigger ID, and return it as an unsigned integer.

6.3.4. GetStandardEvent

This method should extract the sensor data from the `RawDataEvent` input parameter, and fill in the `StandardEvent` by adding the appropriate number of `StandardPlanes` (one per sensor plane). The method signature is:

```
virtual bool GetStandardSubEvent(StandardEvent & out,
                                const Event & in) const;
```

It should return `true` if it successfully updated the `StandardEvent`, or `false` to indicate an error.

6.3.5. GetLCIOEvent

Similar to `GetStandardEvent`, the `GetLCIOEvent` method converts a `RawDataEvent` into a standardized format, in this case LCIO. The signature is:

```
virtual lcio::LCEvent * GetLCIOEvent(const Event * ev) const;
```

It receives the `RawDataEvent` as a parameter, and should return a pointer to a new `LCEvent` if the conversion is successful. In the event of an error, it should return a null pointer.

7. Other Parts of the Framework

The EUDAQ framework contains a number of other parts that may be useful. Those that have not already been described in previous sections will be outlined below.

7.1. FileWriter

The `FileWriter` part of the framework allows different file formats to be written, using a common interface, using a plugin-like system to define new file types. The `FileWriter` class defines the interface that each type must implement, and the `FileWriterFactory` allows code that writes data files to select any available file type, and write it in a generic way, without needing to know details about the particular file format. A number of different file types are already implemented, for a list with descriptions, see page 32

The easiest way to make use of the different `FileWriters`, is to use the `Converter.exe` program (see subsection 4.4.4).

The `FileWriter` base class is defined in the following header:

```
#include "eudaq/FileWriter.hh"
```

In order to implement a new `FileWriter`, a new class must be written, inheriting from the `FileWriter` base class, and implementing the following methods:

```
virtual void StartRun(unsigned);
virtual void WriteEvent(const DetectorEvent &);
virtual uint64_t FileBytes() const;
```

The `StartRun` method is called at the start of each new run with the run number as a parameter. This allows a new file to be opened, and any header information to be written if necessary. Then the `WriteEvent` method is called for each event to be written. Here the `DetectorEvent` can be decoded and processed and the necessary data written to file. The `FileBytes` method should return the number of bytes written to the file. However, it is optional, and may simply return zero if the actual size is not easily known.

7.2. FileReader

Although tools are provided to access the information in the native data files, and to convert them to other formats (such as LCIO, for analysis with the EUTelescope package), in some cases it may be preferable to access the native data directly. For this, the `FileReader` class is provided, allowing a custom program to be written to access a native file and process it as desired.

The constructor takes as an argument the name of the file to be opened, and will read the first event from the file (which should be the BORE). The `NextEvent()` method can then be called to advance through the file. It can optionally take as a parameter the number of events to skip, and will return `true` as long as a new event was read. The currently loaded event can be accessed with the `GetDetectorEvent()` method.

The basic usage is shown below, while a more complete example is available in subsection A.4:

```
#include "eudaq/FileReader.hh"
#include <iostream>

int main(int argc, char ** argv) {
    if (argc != 2) {
        std::cerr << "usage: " << argv[0] << " file" << std::endl;
        return 1;
    }
    eudaq::FileReader reader(argv[1]);
    std::cout << "Opened file: " << reader.Filename() << std::endl;
    std::cout << "BORE:\n" << reader.GetDetectorEvent() << std::endl;
    while (reader.NextEvent()) {
        std::cout << reader.GetDetectorEvent() << std::endl;
    }
    return 0;
}
```

This will open the file specified on the command line, and print out a summary of all the events in there. Be aware that running it as it is may generate a large amount of output, especially with large data files.

7.3. PluginManager

The `PluginManager` handles the different `DataConverterPlugins`, allowing raw data stored in a `RawDataEvent` to be easily converted to a `StandardEvent` or `LCEvent` without having to know the details of all the detector types in there. It is defined in the following header:

```
#include "eudaq/PluginManager.hh"
```

In order to convert the events correctly, the plugins must have access to the information in the BORE. Therefore, before any events may be converted, and for each data file, the `PluginManager` must be initialized as follows:

```
eudaq::PluginManager::Initialize(bore);
```

The `PluginManager` will take care of passing the relevant parts of the BORE to the appropriate `DataConverterPlugins`. The `DetectorEvents` can then be converted as follows:

```
eudaq::StandardEvent sev = eudaq::PluginManager::ConvertToStandard(dev);
```

The `PluginManager` will take care of splitting the `DetectorEvent` into its constituent subevents, and passing them all to the appropriate `DataConverterPlugins` to be inserted into the returned `StandardEvent`. For a slightly more complete example of the `PluginManager`, see the `ExampleReader` in subsection A.4.

7.4. OptionParser

The `OptionParser` is used to simplify parsing of command-line options. It provides a way to specify which arguments a program accepts, with the types, default values and descriptions, so that the help text can be automatically generated, and therefore is always in sync with the code, and all command line programs can have a uniform interface.

All programs using the `OptionParser` will automatically provide a `-h` (and `--help`) option to display the help text, as well as a `-v` (and `--version`) option to display the program version, unless the program explicitly overrides these options with other ones with the same names.

The `OptionParser` is the class that handles the actual parsing of the command line. The signature of the constructor is as follows:

```
OptionParser(const std::string & name, const std::string & version,
             const std::string & desc="", int minargs = -1, int maxargs = -1);
```

The first three arguments are the program name, version and (optionally) description, and these are optionally followed by two numbers specifying the number of arguments expected after the command line options. The default value of -1 for the minimum means no arguments are allowed, and for the maximum means that an arbitrary number may be given (i.e. there is no explicit maximum).

If the automatically generated help text is not sufficient, extra text may also be given to display at the end of the help text, by passing it to the following method:

```
void OptionParser::ExtraHelpText(const std::string & text);
```

This can be used to provide extra information about the options to the program.

Once an `OptionParser` object has been constructed, the different options may be specified. There are two types: `OptionFlag`, which specifies a simple option with no argument, and the template `Option<T>`, which specifies an option taking an argument of type `T`.

The `OptionFlag` constructor has the following signature:

```
OptionFlag(OptionParser & op, const std::string & shortname,
           const std::string & longname, const std::string & desc = "");
```

where `op` is a reference to the `OptionParser` object created previously, that will do the actual parsing of the command line. It then takes two names: a short version (usually a single character) that is used with a single hyphen, and a long version that must be preceded by two hyphens on the command line. Finally, a description may be given that will be displayed in the help text.

The `Option` constructor has the following two signatures, one for normal types, the other for vectors of another type:

```
Option<T>(OptionParser & op, const std::string & shortname,
         const std::string & longname, const T & deflt = T(),
         const std::string & argname = "", const std::string & desc = "");
Option<std::vector<T> >(OptionParser & op, const std::string & shortname,
                       const std::string & longname, const std::string & argname = "",
```

```
const std::string & sep = "", const std::string & desc = "");
```

where, in both cases, the first three arguments are as for `OptionFlag`. The first constructor then takes a default value that will be used in the case the option is not specified on the command line, a name for the argument to the option (to be used in the help text), and a description of the option. The vector version also takes an argument name and a description, but no default value (the default is always an empty vector), instead it takes a separator, which is the string used to separate multiple elements of the vector on the command line. By default (or if an empty string is specified), a comma will be used. Once all the options have been specified, the command line can be parsed, which is done by calling the following method of the `OptionParser` object:

```
OptionParser & OptionParser::Parse(const char ** args);
```

as an argument it takes the list of arguments from the command line (by convention usually called `argv`). If there is an error during parsing, an exception may be thrown; this should be handled by the `HandleMainException` method as described below.

Afterwards the values of the options can be accessed using their `Value()` method. The `IsSet()` method is also available to tell whether an option has been set on the command line (for `OptionFlags` this will hold the same value as the `Value()` method).

Finally, the `OptionParser` has a `HandleMainException` method that provides a way to catch any unhandled exceptions, and either display help if it is a problem with parsing the command line, or otherwise display a standard text informing the user of a problem. It will also catch exceptions of type `MessageException` and display the message, without treating it as an error, so this can be used to exit the program with a message to the user. It is recommended to put the main program inside a `try` block, then call the `HandleMainException` method from a `catch(...)` block, after any other exceptions have been handled (if necessary).

An example use is shown below, illustrating most of what is described above:

```
1 #include "eudaq/OptionParser.hh"
2 #include "eudaq/Utils.hh"
3 #include <iostream>
4
5 int main(int /*argc*/, char ** argv) {
6     eudaq::OptionParser op("Example", "1.0", "An example program", 0);
7     eudaq::OptionFlag test(op, "t", "test", "Enable test");
8     eudaq::Option<double> example(op, "e", "example", 3.14, "value",
9         "Example parameter");
10    eudaq::Option<std::vector<int> > another(op, "a", "another", "values", ";",
11        "Example vector");
12    op.ExtraHelpText("Some more information about this");
13    try {
14        op.Parse(argv);
15        std::cout << "Test: " << (test.IsSet() ? "Enabled\n" : "Disabled\n")
16            << "Example: " << example.Value() << "\n"
```

```

17     << "Another: " << eudaq::to_string(another.Value(), ", ")
18     << std::endl;
19     if (op.NumArgs() == 0) {
20         throw(eudaq::MessageException("No arguments were given"));
21     }
22     for (unsigned i = 0; i < op.NumArgs(); ++i) {
23         std::cout << "Argument " << (i+1) << ": " << op.GetArg(i) << std::endl;
24     }
25 } catch(...) {
26     return op.HandleMainException();
27 }
28 return 0;
29 }

```

Running this program produces the following output:

```

./OptionExample.exe -h
Example version 1.0
An example program

usage: ./OptionExample.exe [options] [0 or more arguments]

options:
  -t --test
      Enable test
  -e --example <value> (default = 42)
      Example parameter
  -a --another <values> (default = )
      Example vector

```

Some more information about this program.

```

./OptionExample.exe
Test: Disabled
Example: 42
Another:
No arguments were given

./OptionExample.exe -t -e 2.718 -a 1;2;3 foo bar
Test: Enabled
Example: 2.718
Another: 1, 2, 3
Argument 1: foo
Argument 2: bar

```

7.5. Timer

The `Timer` class wraps the underlying operating system's timer functions, making them easier to use in a platform independent way. Whenever a `Timer` object is created, it will record the current time. Then at any time in the future, the elapsed time in seconds may be accessed with the `Seconds()` method.

There is also a `Stop()` method to stop the timer counting, so any subsequent calls to `Seconds` will return the same value, and a `Restart()` method to reset the timer's start time to the current time and start counting again. An example use is shown below:

```
#include "eudaq/Timer.hh"

Timer t;
function_a();
cout << "Function A took " << t.Seconds() << " seconds." << endl;
t.Restart();
function_b();
cout << "Function B took " << t.Seconds() << " seconds." << endl;
// wait 3 microseconds
t.Restart();
while (t.Seconds() < 3e-6) {
    // do nothing
}
```

This shows a timer being used to measure the execution time of two functions, and to wait for a small delay. Usually to wait for a delay, it is preferable to use `sleep` (or `mSleep`, see subsection 7.6.4), but in most operating systems the minimum delay for a sleep is around 20 ms (even when using `usleep` which has microsecond resolution) so if the delay must be shorter, a busy loop like above is needed.

7.6. Utils

The `Utils` package is a collection of useful functions and classes too small to merit their own individual files. It is used by including the header:

```
#include "eudaq/Utils.hh"
```

Some of the most useful parts are described here.

7.6.1. to_string

This is a template function that takes (almost) any type and returns the value converted to a string. An optional second argument specifies the minimum number of digits to use (padding with zeroes if necessary).

```
int value = 123;
strfunction(to_string(value));
strfunction(to_string(value, 6));
```

This will pass first the string "123", and then the string "000123" to the function `strfunction`.

7.6.2. `from_string`

This template function is the inverse of `to_string`. It takes as arguments a string and a default value of type T, and returns an object of type T initialised from the string. If it is not possible to convert the string to the required type, the default value is returned instead.

```
std::string value = "456";
intfunction(from_string(value, 0));
```

This will call `intfunction` with the integer value 456.

7.6.3. `hexdec`

This is a class to facilitate printing numbers in both hexadecimal and decimal. It is used similarly to `to_string`, but when printed, it will display the value in hexadecimal, followed by the value in decimal in parentheses. The hexadecimal values will be padded to the full width of the type, unless a second argument is given specifying the minimum number of hex digits to display.

```
short value = 789;
cout << hexdec(value) << endl
     << hexdec(value, 0) << endl;
```

This will display:

```
0x0315 (789)
0x315 (789)
```

If the result is required in a string, instead of being printed, this can be achieved with `to_string(hexdec(value))`.

7.6.4. `mSleep`

This is a wrapper around the operating system's `sleep/usleep` (or equivalent) function. It takes as an argument the number of milliseconds to sleep. The advantage of this function is that it will work on Linux, Mac OS X and Windows, as it will automatically call the correct underlying function.

8. Reporting Issues

The GitHub server, on which EUDAQ is hosted, provides a system for reporting bugs and for requesting new features. It is accessible at the following address:

<https://github.com/eudaq/eudaq/issues>.

Here you may submit new reports (you are required to register first to do this), or follow the status of existing bugs and feature requests. This is recommended over (or at least, as well as) sending an email to the developers, as it ensures a record of the issue is available, and others may follow the progress.

9. Developing and Contributing to EUDAQ

9.1. Regression Testing

If a CMake version later than 2.8.0 is found and Python is installed together with the `numpy` package, several regression tests are made available that can be executed through CTest. The tests are based on the Python wrapper around EUDAQ components as described in section 4.2.10. Run the tests by typing

```
cd build
cmake ..
ctest
```

This starts the script `etc/tests/run_dummydataproduction.py` which runs a short DAQ session with instances of RunControl, DataCollector and a (dummy) Producer and compares the output to a reference file stored on AFS at DESY. If your system is set up correctly, you have access to the reference file, and the basic components of the EUDAQ library work, all tests should pass. To see the output of failing tests, you can add the `--output-on-failure` parameter to the CTest command.

These basic tests can easily be extended to test other parts of the core framework or of your own producer. Take a look at the `etc/tests/testing.cmake` CMake script and the central `CMakeLists.txt` file where it is included for an example of how to set up tests with CTest.

The automated nightly tests are defined in CMake scripts located in `etc/tests/nightly` and are executed by the scripts `run_nightly.sh` and `run_nightly.bat` for Unix and Windows platforms, respectively. In addition to the dummy run described above, the nightly tests check out all changes from the central repository, build the full code base, and submit all results to the CDash webserver hosted at DESY: <http://aidasoft.desy.de/CDash/index.php?project=EUDAQ>

9.2. Committing Code to the Main Repository

If you would like to contribute your code back into the main repository, please follow the “fork & pull request” strategy:

- Create a user account on github, log in
- “Fork” the (main) project on github (using the button on the page of the main repo)
- *Either*: clone from the newly forked project and add ‘upstream’ repository to local clone (change user names in URLs accordingly):

```
git clone https://github.com/hperrey/eudaq eudaq
cd eudaq
git remote add upstream https://github.com/eudaq/eudaq.git
```

- *or* if edits were made to a previous checkout of upstream: rename origin to upstream, add fork as new origin:

```
cd eudaq
git remote rename origin upstream
git remote add origin https://github.com/hperrey/eudaq
git remote -v show
```

- Optional: edit away on your local clone! But keep in sync with the development in the upstream repository by running

```
git fetch upstream      # download named heads or tags
git pull upstream master # merge changes into your branch
```

on a regular basis. Replace `master` by the appropriate branch if you work on a separate one. Don't forget that you can refer to issues in the main repository anytime by using the string `eudaq/eudaq#XX` in your commit messages, where `XX` stands for the issue number, e.g.

```
[...]. this addresses issue eudaq/eudaq#1
```

- Push the edits to origin (our fork)

```
git push origin
```

(defaults to `git push origin master` where origin is the repo and master the branch)

- Verify that your changes made it to your github fork and then click there on the “compare & pull request” button
- Summarize your changes and click on “send”
- Thank you!

Working together on a branch: If you have a copy installed, and want to update it to the latest version, you do not need to clone the repository again, just change to the `eudaq` directory use the command:

```
git pull
```

to update your local copy with all changes committed to the central repository.

A. Source Code

This section contains example code to illustrate the concepts in the manual, when they are too long to be included in the main section.

All files are also present in the EUDAQ distribution; so if possible those versions should be used, since they may be more up to date than the manual.

A.1. Example Config File

Latest version available at:

<https://github.com/eudaq/eudaq/blob/master/conf/ExampleConfig.conf>

```
1 # This is an example config file, you can adapt it to your needs.
2 # All text following a # character is treated as comments
3
4 [RunControl]
5 RunSizeLimit = 1000000000
6
7 [DataCollector]
8 FilePattern = "../data/run$6R$X"
9
10 [LogCollector]
11 SaveLevel = EXTRA
12 PrintLevel = INFO
13
14 [Producer.Example]
15 ConfParameter = 123
16
17 [Producer.TLU]
18 AndMask = 0xf
19 OrMask = 0
20 VetoMask = 0
21 DutMask = 20
22 TriggerInterval = 0
23 TrigRollover = 0
24 #DUTInput3=LEMO
```

A.2. Example Producer

Latest version available at:

<https://github.com/eudaq/eudaq/blob/master/main/exe/src/ExampleProducer.cxx>

```

1 #include "eudaq/Configuration.hh"
2 #include "eudaq/Producer.hh"
3 #include "eudaq/Logger.hh"
4 #include "eudaq/RawDataEvent.hh"
5 #include "eudaq/Timer.hh"
6 #include "eudaq/Utils.hh"
7 #include "eudaq/OptionParser.hh"
8 #include "eudaq/ExampleHardware.hh"
9 #include <iostream>
10 #include <ostream>
11 #include <vector>
12
13 #include "eudaq/Status.hh"
14
15 // A name to identify the raw data format of the events generated
16 // Modify this to something appropriate for your producer.
17 static const std::string EVENT_TYPE = "Example";
18
19 // Declare a new class that inherits from eudaq::Producer
20 class ExampleProducer : public eudaq::Producer {
21     public:
22
23     // The constructor must call the eudaq::Producer constructor with the name
24     // and the runcontrol connection string, and initialize any member ←
25     // variables.
26     ExampleProducer(const std::string & name, const std::string & runcontrol)
27         : eudaq::Producer(name, runcontrol),
28         m_run(0), m_ev(0), stopping(false), done(false) {}
29
30     // This gets called whenever the DAQ is initialised
31     virtual void OnInitialise(const eudaq::Configuration & init) {
32         try {
33             std::cout << "Reading: " << init.Name() << std::endl;
34
35             // Do any initialisation of the hardware here
36             // "start-up configuration", which is usually done only once in the ←
37             // beginning
38             // Configuration file values are accessible as config.Get(name, ←
39             // default)
40             m_exampleInitParam = init.Get("InitParameter", 0);
41
42             // send information

```

```

40     // Message as cout in the terminal of your producer
41     std::cout << "Initialise with parameter = " << m_exampleInitParam << ←
        std::endl;
42     // or to the LogCollector, depending which log level you want. These ←
        are the possibilities just as an example here:
43     EUDAQ_INFO("Initialise with parameter = " + m_exampleInitParam);
44     //EUDAQ_DEBUG("Debug Message to the LogCollector from ←
        ExampleProducer");
45     //EUDAQ_EXTRA("Extra Message to the LogCollector from ←
        ExampleProducer");
46     //EUDAQ_INFO("Info Message to the LogCollector from ExampleProducer");
47     //EUDAQ_WARN("Warn Message to the LogCollector from ExampleProducer");
48     //EUDAQ_ERROR("Error Message to the LogCollector from ←
        ExampleProducer");
49     //EUDAQ_USER("User Message to the LogCollector from ExampleProducer");
50
51     // EUDAQ_THROW throws an error, thus here goes to catch! With that ←
        you can simulate errors...
52     //EUDAQ_THROW("User Message to the LogCollector from ExampleProducer");
53
54     // send it to your hardware
55     hardware.Setup(m_exampleInitParam);
56
57     // At the end, set the ConnectionState that will be displayed in the ←
        Run Control.
58     // and set the state of the machine.
59     SetConnectionState(eudaq::ConnectionState::STATE_UNCONF, ←
        "Initialised (" + init.Name() + ")");
60 }
61 catch (...) {
62     // Message as cout in the terminal of your producer
63     std::cout << "Unknown exception" << std::endl;
64     // Message to the LogCollector
65     EUDAQ_ERROR("Error Message to the LogCollector from ExampleProducer");
66     // Otherwise, the State is set to ERROR
67     SetConnectionState(eudaq::ConnectionState::STATE_ERROR, ←
        "Initialisation Error");
68 }
69 }
70
71 // This gets called whenever the DAQ is configured
72 virtual void OnConfigure(const eudaq::Configuration & config) {
73     try {
74         std::cout << "Reading: " << config.Name() << std::endl;
75
76         // Do any configuration of the hardware here

```

```

77     // Configuration file values are accessible as config.Get(name, ←
        default)
78     m_exampleConfParam = config.Get("ConfParameter", 0);
79     std::cout << "Example Configuration Parameter = " << ←
        m_exampleConfParam << std::endl;
80     hardware.Setup(m_exampleConfParam);
81
82     // At the end, set the ConnectionState that will be displayed in the ←
        Run Control.
83     // and set the state of the machine.
84     SetConnectionState(eudaq::ConnectionState::STATE_CONF, "Configured ←
        (" + config.Name() + ")");
85 }
86 catch (...) {
87     // Otherwise, the State is set to ERROR
88     printf("Unknown exception\n");
89     SetConnectionState(eudaq::ConnectionState::STATE_ERROR, ←
        "Configuration Error");
90 }
91 }
92
93 // This gets called whenever a new run is started
94 // It receives the new run number as a parameter
95 // And sets the event number to 0 (internally)
96 virtual void OnStartRun(unsigned param) {
97     try {
98
99         m_run = param;
100        m_ev = 0;
101
102        std::cout << "Start Run: " << m_run << std::endl;
103
104        // It must send a BORE (Begin-Of-Run Event) to the Data Collector
105        eudaq::RawDataEvent bore(eudaq::RawDataEvent::BORE(EVENT_TYPE, m_run));
106        // You can set tags on the BORE that will be saved in the data file
107        // and can be used later to help decoding
108        bore.SetTag("EXAMPLE", eudaq::to_string(m_exampleConfParam));
109        // Starting your hardware
110        hardware.PrepareForRun();
111        // Send the event to the Data Collector
112        SendEvent(bore);
113
114        // At the end, set the ConnectionState that will be displayed in the ←
        Run Control.
115        SetConnectionState(eudaq::ConnectionState::STATE_RUNNING, "Running");
116    }

```

```

117     catch (...) {
118         // Otherwise, the State is set to ERROR
119         printf("Unknown exception\n");
120         SetConnectionState(eudaq::ConnectionState::STATE_ERROR, "Starting ←
           Error");
121     }
122 }
123
124 // This gets called whenever a run is stopped
125 virtual void OnStopRun() {
126     try {
127         // Set a flag to signal to the polling loop that the run is over and ←
           it is in the stopping process
128         stopping = true;
129
130         // wait until all events have been read out from the hardware
131         while (stopping) {
132             eudaq::mSleep(20);
133             //std::cout<<"Does hardware have pending? ←
           "<<hardware.EventsPending()<<"\n";
134         }
135         // Send an EORE after all the real events have been sent
136         // You can also set tags on it (as with the BORE) if necessary
137         SendEvent(eudaq::RawDataEvent::EORE("Test", m_run, ++m_ev));
138
139         // At the end, set the ConnectionState that will be displayed in the ←
           Run Control.
140         // Due to the definition of FSM, it should go to STATE_CONF.
141         if (m_connectionstate.GetState() != ←
           eudaq::ConnectionState::STATE_ERROR)
142             SetConnectionState(eudaq::ConnectionState::STATE_CONF);
143     }
144     catch (...) {
145         // Otherwise, the State is set to ERROR
146         printf("Unknown exception\n");
147         SetConnectionState(eudaq::ConnectionState::STATE_ERROR, "Stopping ←
           Error");
148     }
149 }
150
151 // This gets called when the Run Control is terminating,
152 // we should also exit.
153 virtual void OnTerminate() {
154     std::cout << "Terminating..." << std::endl;
155     done = true;
156 }

```

```

157
158 // This loop is running in the main
159 // This is just an example, adapt it to your hardware
160 void ReadoutLoop() {
161     try {
162         // Loop until Run Control tells us to terminate using the done flag
163         while (!done) {
164             if (!hardware.EventsPending()) {
165                 // No events are pending, so check if the run is stopping
166                 if (stopping) {
167                     // if so, signal that there are no events left
168                     stopping = false;
169                 }
170                 // Now sleep for a bit, to prevent chewing up all the CPU
171                 eudaq::mSleep(20);
172                 // Then restart the loop
173                 continue;
174             }
175             // If the Producer is not in STATE_RUNNING, it will restart the loop
176             if (GetConnectionState() != eudaq::ConnectionState::STATE_RUNNING) {
177                 // Now sleep for a bit, to prevent chewing up all the CPU
178                 eudaq::mSleep(20);
179                 // Then restart the loop
180                 continue;
181             }
182
183             // If we get here, there must be data to read out
184             // Create a RawDataEvent to contain the event data to be sent
185             eudaq::RawDataEvent ev(EVENT_TYPE, m_run, m_ev);
186
187             for (unsigned plane = 0; plane < hardware.NumSensors(); ++plane) {
188                 // Read out a block of raw data from the hardware
189                 std::vector<unsigned char> buffer = hardware.ReadSensor(plane);
190                 // Each data block has an ID that is used for ordering the ←
191                 // planes later
192                 // If there are multiple sensors, they should be numbered ←
193                 // incrementally
194
195                 // Add the block of raw data to the event
196                 ev.AddBlock(plane, buffer);
197             }
198             hardware.CompletedEvent();
199             // Send the event to the Data Collector
200             SendEvent(ev);
201             // Now increment the event number
202             m_ev++;

```



```

201     }
202 }
203 catch (...) {
204     // Otherwise, the State is set to ERROR
205     printf("Unknown exception\n");
206     SetConnectionState(eudaq::ConnectionState::STATE_ERROR, "Error ←
        during running");
207 }
208 }
209
210 private:
211     // This is just a dummy class representing the hardware
212     // It here basically that the example code will compile
213     // but it also generates example raw data to help illustrate the decoder
214     eudaq::ExampleHardware hardware;
215     unsigned m_run, m_ev, m_exampleConfParam, m_exampleInitParam;
216     bool stopping, done;
217 };
218
219 // The main function that will create a Producer instance and run it
220 int main(int /*argc*/, const char ** argv) {
221     // You can use the OptionParser to get command-line arguments
222     // then they will automatically be described in the help (-h) option
223     eudaq::OptionParser op("EUDAQ Example Producer", "1.0",
224         "Just an example, modify it to suit your own needs");
225     eudaq::Option<std::string> rctrl(op, "r", "runcontrol",
226         "tcp://localhost:44000", "address",
227         "The address of the RunControl.");
228     eudaq::Option<std::string> level(op, "l", "log-level", "NONE", "level",
229         "The minimum level for displaying log messages locally");
230     eudaq::Option<std::string> name(op, "n", "name", "Example", "string",
231         "The name of this Producer");
232     try {
233         // This will look through the command-line arguments and set the options
234         op.Parse(argv);
235         // Set the Log level for displaying messages based on command-line
236         EUDAQ_LOG_LEVEL(level.Value());
237         // Create a producer
238         ExampleProducer producer(name.Value(), rctrl.Value());
239         // And set it running...
240         producer.ReadoutLoop();
241         // When the readout loop terminates, it is time to go
242         std::cout << "Quitting" << std::endl;
243     } catch (...) {
244         // This does some basic error handling of common exceptions
245         return op.HandleMainException();

```

```
246 }  
247 return 0;  
248 }
```

A.3. Example DataConverterPlugin

Latest version available at:

<https://github.com/eudaq/eudaq/blob/master/main/lib/plugins/ExampleConverterPlugin.cc>

```

1 #include "eudaq/DataConverterPlugin.hh"
2 #include "eudaq/StandardEvent.hh"
3 #include "eudaq/Utils.hh"
4
5 // All LCIO-specific parts are put in conditional compilation blocks
6 // so that the other parts may still be used if LCIO is not available.
7 #if USE_LCIO
8 #include "IMPL/LCEventImpl.h"
9 #include "IMPL/TrackerRawDataImpl.h"
10 #include "IMPL/LCCollectionVec.h"
11 #include "lcio.h"
12 #endif
13
14 namespace eudaq {
15
16     // The event type for which this converter plugin will be registered
17     // Modify this to match your actual event type (from the Producer)
18     static const char *EVENT_TYPE = "Example";
19
20     // Declare a new class that inherits from DataConverterPlugin
21     class ExampleConverterPlugin : public DataConverterPlugin {
22
23     public:
24         // This is called once at the beginning of each run.
25         // You may extract information from the BORE and/or configuration
26         // and store it in member variables to use during the decoding later.
27         virtual void Initialize(const Event &bore, const Configuration &cnf) {
28             m_exampleparam = bore.GetTag("EXAMPLE", 0);
29 #ifndef WIN32 // some linux Stuff //$$change
30             (void)cnf; // just to suppress a warning about unused parameter cnf
31 #endif
32         }
33
34         // This should return the trigger ID (as provided by the TLU)
35         // if it was read out, otherwise it can either return (unsigned)-1,
36         // or be left undefined as there is already a default version.
37         virtual unsigned GetTriggerID(const Event &ev) const {
38             static const unsigned TRIGGER_OFFSET = 8;
39             // Make sure the event is of class RawDataEvent
40             if (const RawDataEvent *rev = dynamic_cast<const RawDataEvent *>(&ev)) {
41                 // This is just an example, modified it to suit your raw data format
42                 // Make sure we have at least one block of data, and it is large enough

```

```

43     if (rev->NumBlocks() > 0 &&
44         rev->GetBlock(0).size() >= (TRIGGER_OFFSET + sizeof(short))) {
45         // Read a little-endian unsigned short from offset TRIGGER_OFFSET
46         return getlittleendian<unsigned short>(
47             &rev->GetBlock(0)[TRIGGER_OFFSET]);
48     }
49 }
50 // If we are unable to extract the Trigger ID, signal with (unsigned)-1
51 return (unsigned)-1;
52 }
53
54 // Here, the data from the RawDataEvent is extracted into a StandardEvent.
55 // The return value indicates whether the conversion was successful.
56 // Again, this is just an example, adapted it for the actual data layout.
57 virtual bool GetStandardSubEvent(StandardEvent &sev,
58                                 const Event &ev) const {
59     // If the event type is used for different sensors
60     // they can be differentiated here
61     std::string sensortype = "example";
62     // Create a StandardPlane representing one sensor plane
63     int id = 0;
64     StandardPlane plane(id, EVENT_TYPE, sensortype);
65     // Set the number of pixels
66     int width = 100, height = 50;
67     plane.SetSizeRaw(width, height);
68     // Set the trigger ID
69     plane.SetTLUEvent(GetTriggerID(ev));
70     // Add the plane to the StandardEvent
71     sev.AddPlane(plane);
72     // Indicate that data was successfully converted
73     return true;
74 }
75
76 #if USE_LCIO
77 // This is where the conversion to LCIO is done
78 virtual lcio::LCEvent *GetLCIOEvent(const Event * /*ev*/) const {
79     return 0;
80 }
81 #endif
82
83 private:
84 // The constructor can be private, only one static instance is created
85 // The DataConverterPlugin constructor must be passed the event type
86 // in order to register this converter for the corresponding conversions
87 // Member variables should also be initialized to default values here.
88 ExampleConverterPlugin()

```

```
89     : DataConverterPlugin(EVENT_TYPE), m_exampleparam(0) {}
90
91     // Information extracted in Initialize() can be stored here:
92     unsigned m_exampleparam;
93
94     // The single instance of this converter plugin
95     static ExampleConverterPlugin m_instance;
96 }; // class ExampleConverterPlugin
97
98 // Instantiate the converter plugin instance
99 ExampleConverterPlugin ExampleConverterPlugin::m_instance;
100
101 } // namespace eudaq
```

A.4. Example Reader

Latest version available at:

<https://github.com/eudaq/eudaq/blob/master/main/exe/src/ExampleReader.cxx>

```

1 #include "eudaq/FileReader.hh"
2 #include "eudaq/PluginManager.hh"
3 #include "eudaq/OptionParser.hh"
4 #include <iostream>
5
6 static const std::string EVENT_TYPE = "Example";
7
8 int main(int /*argc*/, const char ** argv) {
9     // You can use the OptionParser to get command-line arguments
10    // then they will automatically be described in the help (-h) option
11    eudaq::OptionParser op("EUDAQ Example File Reader", "1.0",
12        "Just an example, modify it to suit your own needs",
13        1);
14    eudaq::OptionFlag doraw(op, "r", "raw", "Display raw data from events");
15    eudaq::OptionFlag docon(op, "c", "converted", "Display converted events");
16    try {
17        // This will look through the command-line arguments and set the options
18        op.Parse(argv);
19
20        // Loop over all filenames
21        for (size_t i = 0; i < op.NumArgs(); ++i) {
22
23            // Create a reader for this file
24            eudaq::FileReader reader(op.GetArg(i));
25
26            // Display the actual filename (argument could have been a run number)
27            std::cout << "Opened file: " << reader.Filename() << std::endl;
28
29            // The BORE is now accessible in reader.GetDetectorEvent()
30            if (docon.IsSet()) {
31                // The PluginManager should be initialized with the BORE
32                eudaq::PluginManager::Initialize(reader.GetDetectorEvent());
33            }
34
35            // Now loop over all events in the file
36            while (reader.NextEvent()) {
37                if (reader.GetDetectorEvent().IsEORE()) {
38                    std::cout << "End of run detected" << std::endl;
39                    // Don't try to process if it is an EORE
40                    break;
41                }
42

```

```
43     if (doraw.IsSet()) {
44         // Display summary of raw event
45         //std::cout << reader.GetDetectorEvent() << std::endl;
46
47         try {
48             // Look for a specific RawDataEvent, will throw an exception if ←
49             // not found
50             const eudaq::RawDataEvent & rev =
51                 reader.GetDetectorEvent().GetRawSubEvent(EVENT_TYPE);
52             // Display summary of the Example RawDataEvent
53             std::cout << rev << std::endl;
54         } catch (const eudaq::Exception & ) {
55             std::cout << "No " << EVENT_TYPE << " subevent in event "
56                 << reader.GetDetectorEvent().GetEventNumber()
57                 << std::endl;
58         }
59
60         if (docon.IsSet()) {
61             // Convert the RawDataEvent into a StandardEvent
62             eudaq::StandardEvent sev =
63                 eudaq::PluginManager::ConvertToStandard(reader.GetDetectorEvent());
64
65             // Display summary of converted event
66             std::cout << sev << std::endl;
67         }
68     }
69 }
70
71 } catch (...) {
72     // This does some basic error handling of common exceptions
73     return op.HandleMainException();
74 }
75 return 0;
76 }
```

B. Introduction to the build system and project files on Windows

B.1. MSBUILD

This is the program that processes the project (solution) files and feeds it to the compiler and linker. If you have a working project file it is more or less straight forward. It has a very simple syntax:

```
MSBUILD.exe MyApp.sln /t:Rebuild /p:Configuration=Release
```

myApp.sln is the file you want to Process. The parameter `/target` (short `/t`) tells msbuild what to do in this case rebuild. You have all the options you need like: clean, build and rebuild. You can also specify your own targets. With the “parameter property” switch you can change the properties of your Project. Let’s say you want to compile EUDAQ, you go in the build folder where the solution (sln) file is and type:

```
MSBUILD.exe EUDAQ.sln /p:Configuration=Release
```

One thing one has to keep in mind is that there are some default configurations. The default is a debug build for x86. If you want to have it different then you need to specify it in the command line. And one thing you want to have is a release build! With the `/p` switch you can overwrite properties like in this case the configuration. But you could also overwrite the compiler version it should use. Let’s say you want to use VS 2013 then you have to specify it by writing:

```
MSBUILD.exe EUDAQ.sln /p:PlatformToolset=v120 /p:Configuration=Release
```

But be careful when changing the compiler settings. It is possible that some then link against an incompatible version of your external libraries.

B.2. Project Files

Project files are the Visual Studio equivalent to Makefiles. The Project files have a very easy syntax but a complicated mechanism behind it. Making changes to an existing file is very easy. Writing a new one from scratch is expert level. But also, in most cases, pointless because CMake does it for you. Therefore usually one gets a finished Project file that was auto created by CMake and one just wants to make some minor changes to it, therefore it is enough to know where one can tweak around.

Please remember to adjust the CMake files when you are done accordingly, so that your changes are reproduced and not overwritten on the next CMake run.

Let’s start easy and assume you want to change the output directory. You can do this by adding the following line to the corresponding Property group.

```
<PropertyGroup ←  
  Condition=" '$(Configuration)|$(Platform)'=='Release|Win32' ">  
    <OutDir>..\..\Windows Binaries\</OutDir>
```



```
</PropertyGroup>
```

Or let's say you want to change the compiler version. You can do this by changing the platform toolset to the version you need. You can find this option in

```
<PropertyGroup ←
  Condition=" '$(Configuration)|$(Platform)'=='Debug|Win32'" ←
  Label=" Configuration">
  ...
<PlatformToolset>v110</PlatformToolset>
</PropertyGroup>
```

V110 stands for Visual Studio 2012. V120 stands for VS 2013 and so on. The next interesting switches are in here:

```
<ItemDefinitionGroup Condition=" '$(Configuration)|$(Platform)'=='Debug|Win32'">
  <ClCompile>
    <PrecompiledHeader></PrecompiledHeader>
    <WarningLevel>Level3</WarningLevel>
    <Optimization>Disabled</Optimization>
    <PreprocessorDefinitions>
      WIN32;
      DEBUG;
      CONSOLE;
      %(PreprocessorDefinitions)
    </PreprocessorDefinitions>
    <AdditionalIncludeDirectories>
      ..\..\main\include;
      ..\..\extern\pthread-win32\include;
      ..\..\tlu\include;
      ..\..\extern\ZestSC1\windows 7\Inc;
      ..\..\extern\libusb-win32-bin-1.2.6.0\include
    </AdditionalIncludeDirectories>
  </ClCompile>
  <Link>
    <SubSystem>Console</SubSystem>
    <GenerateDebugInformation>true</GenerateDebugInformation>
    <AdditionalLibraryDirectories>
      ..\..\extern\libusb-win32-bin-1.2.6.0\lib\msvc\;
      ..\..\extern\ZestSC1\windows 7\lib\x86\
    </AdditionalLibraryDirectories>
    <AdditionalDependencies>
      ZestSC1.lib;libusb.lib;kernel32.lib;user32.lib;gdi32.lib;winspool.lib;comdlg32.lib;
      advapi32.lib;shell32.lib;ole32.lib;oleaut32.lib;uuid.lib;odbc32.lib;
      odbccp32.lib;%(AdditionalDependencies)
    </AdditionalDependencies>
  </Link>
</ItemDefinitionGroup>
```

An Item definition Group is the place where you define your items. One can compare Items to a struct in C++; it is an object that contains different types of information. The Condition statement works like an IF in C++.

In this article you can find all the possibilities you have: <http://msdn.microsoft.com/de-de/library/7szfhaft.aspx> In the next line you are defining an item called “CLCompile” and you give it the some attributes like “PreprocessorDefinitions” or “AdditionalIncludeDirectories”. This Object contains all the information that gets sent to the compiler. That means all the compiler flags are set here. The actual files are included later in the project file. So for now you have only defined how you want to compile your files but not what files you want to compile. AdditionalIncludeDirectories does exactly what you think it does. It understands all relative paths and path with environment variables exactly as it should. Next thing is “PreprocessorDefinitions”. It also works exactly as you think it does. That means you can either define just names for your #ifdef statements in the code or you can define macros like

```
<PreprocessorDefinitions>
    SOMEVALUE=3;
    WIN32;
    _DEBUG;
    _CONSOLE;
    %(PreprocessorDefinitions)
</PreprocessorDefinitions>
```

Then you can call in your code SOMEVALUE and it will be 3. I do not know if it is possible to define macro function like

```
#define x_square(x) x*x.
```

```
<AdditionalDependencies>
    $(myFancyLibPath)\*.lib ;
    odbccp32.lib;%(AdditionalDependencies)
</AdditionalDependencies>
```

And it will link against all *.lib files in this directory.

Next thing you need to know is where to put your files you want to compile. Somewhere below the ItemDefinitionGroup there is an ItemGroup which contains the Include statements. It looks like this:

```
<ItemGroup>
    <ClCompile Include="src\someFile.cc" />
    <ClCompile Include="src\someOtherFile.cc" />
    ...
    <ClCompile Include="src\*.cpp" />
    ...
</ItemGroup>
```

Here you can either put individual files or groups of files in. But be careful that you don't include the same file twice. There is also an ItemGroup which contains the include files. This one seems to be more important for the IDE of VS so that it shows the header files in the Solution Explorer.

A typical use case is that you wrote your own *Data Converter Plugin*. This file needs to be mentioned here!

What you won't find in the project file is the section that passes the files to the compiler. This part is hidden behind the following import statement:

```
<Import Project="$(VCTargetsPath)\Microsoft.Cpp.targets" />
```

It is usually not required to modify this file. But if you want to view it you can find it in this folder:

```
C:\Program Files (x86)\MSBuild\Microsoft.Cpp\v4.0\
```

This file is written neither to be very clear nor understandable, so better check out the documentation pages such as:

<http://msdn.microsoft.com/en-us/library/dd293626.aspx>

B.3. Known Problems

- The environment variables are pulled in as properties therefore they can be overwritten in the project file or in the "vcxproj.user" file. So if for example your QT Project won't compile and keeps complaining about not finding the correct directory make sure you are not overwriting the QTDIR environment Variable with a Property.

C. Online Monitor Configuration Settings

C.1. Configuration Sections Overview

we have the following Section Keywords, to be put in [].

- [General]
- [Correlations]
- [Clusterizer]
- [HotPixelFinder]
- [Mimosa26]

C.2. Configuration options in [General]

SnapshotDir *string*

Stores the location of snapshots from the online monitor

SnapshotFormat *string*

Which Format to use for the snapshots, e.g. ".pdf"

C.3. Configuration options in [Correlations]

MinClusterSize *int*

Which minimum cluster size to use for the correlation plots

DisablePlanes *int,int,int*

List of planes to disable, separates by a ","

C.4. Configuration options in [Clusterizer]

C.5. Configuration options in [HotPixelFinder]

HotPixelCut *float*

Cut above which a pixel is considered "hot"

C.6. Configuration options in [Mimosa26]

Mimosa26_max_sections *int*

Number of section of the Mimosa 26 chip, default is 4

Mimosa26_section_boundary *int*

Number of pixels in a Mimosa26 section, default is 288

C.7. Configuration Example

[General]

SnapShotDir = "/scratch/eudet/EUDAQ/bin/"

SnapShotFormat = ".pdf"

[Correlations]

MinClusterSize = 2

DisablePlanes = 2,3

[Clusterizer]

[HotPixelFinder]

HotPixelCut = 0.05

[Mimosa26]

Mimosa26_max_sections = 4

Mimosa26_section_boundary = 288

Glossary

BORE beginning-of-run-event, basically a run header.

CDS correlated double sampling, when two frames are acquired, one before and one after the trigger, and then subtracted to get the actual signal.

DUT device under test.

EORE end-of-run-event, basically a run trailer.

EUDRB teawt.

FSM finite-state machine.

LCIO Linear Collider I/O, the file format used by the analysis software.

NI the National Instrument system, for reading out the Mimosa 26 sensors.

TLU the Trigger Logic Unit.

Acknowledgements

This project receives funding from the European Unions Horizon 2020 Research and Innovation programme under Grant Agreement no. 654168. The support is gratefully acknowledged. *Disclaimer:* The information herein only reflects the views of its authors and not those of the European Commission and no warranty expressed or implied is made with regard to such information or its use.

Before, this work was supported by the Commission of the European Communities under the 6th Framework Programme “Structuring the European Research Area,” contract number RII3-026126, and received funding from the European Commission under the FP7 Research Infrastructures project AIDA, grant agreement no. 262025.

References

- [1] J. Dreyling-Eschweiler and H. Jansen, “EUDET-type beam telescopes”, *Online Wiki*.
URL <https://telescopes.desy.de/>
- [2] P. Roloff, “The EUDET high resolution pixel telescope”, *Nucl. Instrum. Meth.*, **A604**, (2009), 265–268.
- [3] H. Jansen, S. Spannagel, J. Behr, A. Bulgheroni, G. Claus *et al.*, “Performance of the EUDET-type beam telescopes”, *EPJ Techniques and Instrumentation*, **3** (1), (2016), 7.
URL <http://dx.doi.org/10.1140/epjti/s40485-016-0033-2>
- [4] A. Bulgheroni, “EUTelescope, the JRA1 tracking and reconstruction software: a status report”, *EUDET-Memo-2008-48*.
URL <http://www.eudet.org/e26/e28/e615/e835/eudet-memo-2008-48.pdf>

- [5] S. Spannagel, “Test Beam Measurements for the Upgrade of the CMS Pixel Detector and Measurement of the Top Quark Mass from Differential Cross Sections”, Ph.D. thesis, U. Hamburg, Dept. Phys., Hamburg (2016).
URL <http://bib-pubdb1.desy.de/search?cc=Publication+Database&of=hd&p=reportnumber:DESY-THESIS-2016-010>
- [6] GitHub, “Mastering Markdown”, *GitHub*.
URL <https://guides.github.com/features/mastering-markdown/>
- [7] D. G. Cussans, “Description of the JRA1 Trigger Logic Unit (TLU), v0.2c”, *EUDET-Memo-2009-04*.
URL <http://www.eudet.org/e26/e28/e42441/e57298/EUDET-MEMO-2009-04.pdf>
- [8] Git, “Git – local-branching-on-the-cheap”, *Online Article*.
URL <https://git-scm.com/>
- [9] E. Developers, “EUDAQ code on GutHub”, *GitHub*.
URL <https://github.com/eudaq/eudaq>
- [10] J. Dreyling-Eschweiler and H. Jansen, “EUDET-type beam telescopes”, *Online Wiki - User Manual*.
URL https://telescopes.desy.de/User_manual
- [11] D. Shirokova, “Software Development for a common DAQ at test”, *DESY Summer-studetnts*.
- [12] E. Configurations, “EUDAQ configurations files and scripts on GitHub”, *GitHub*.
URL <https://github.com/eudaq/eudaq-configuration>