

ATL-COM-GEN-2009-013

AtlasDcsWebViewer: a web access to the Atlas DCS data

Author: Urban Bitenc, Saverio D'Auria, Pamela Ferrari, Luca Luisa
Status: Work in Progress
Keywords: DCS on-line database

Abstract

The DCS data from the various Atlas subdetectors are stored to the Oracle database using the RDB manager in PVSS. All subdetectors use either the same or very similar schema. It is clear the need to retrieve these data, for a variety of purposes. The effort coordinated within the Inner Detector has produced a web-based tool to search the Atlas DCS Oracle database and to display the results. This tool has been easily extended to acces the data from other Atlas subdetectors. In this note we describe the structure of the AtlasDcsWebViewer and its use.



Contents

1	Introduction	3
2	User Manual	4
2.1	Supported browsers	4
2.2	URL	4
2.3	Choosing the Detector Area	4
2.4	Choosing the Time Interval	5
2.5	Selecting One Data Point Element	5
2.6	The Display	5
2.7	BitAND in Display	6
2.8	Making a list of Data Point Elements	6
2.9	Using the Free-Space Filter	6
2.10	Keeping and sharing the lists	7
2.11	Queries on values	7
2.11.1	Out-of-range queries	8
2.11.2	In-range queries	8
2.11.3	Step queries	8
2.11.4	BitAND in queries on values	9
2.12	Special Queries	9
2.13	Histograms	9
2.14	Results	10
2.15	Saving the results (numeric)	10
2.16	Saving the results (graphic)	10
2.17	Configuration and Personal Settings	11
2.18	Troubleshooting	11
3	Architecture	11
	References	15

1 Introduction

The Atlas detector has a very large number of “channels”, which are controlled by the SCADA system “PVSS” [2] in a distributed system, which is divided in subdetectors and within each subdetector individual projects (often corresponding to PC’s) are run to control a part of the apparatus. The Detector Control System controls and monitors the multitude of parameters of each subdetector and records the data to the on-line Oracle database. The data are also available for display trends and for making queries to the same PVSS system. However during the commissioning phase it clearly emerged the need of accessing the DCS data from other environments both for display and for further processing. The primary reason for accessing the DCS database with queries “on demand” and not with standardized plots is the fast resolution of possible problems, to investigate an event that happened at a given time (e.g. one channel tripping), to correlate events that happened in different parts of the apparatus, or even different projects, and to query the database to search for such events, even if they went un noticed to the primary on-line controls (PVSS).

A set of requirements for the Inner Detector was listed in a series of meetings. These are:

- web-based: the plots and results should be available in a web browser, without installing dedicated software.
- possibility to make queries on values: for example, retrieve all channels which from time A to time B, had a current larger than 20 mA and/or a Temperature higher than 37 degrees.
- possibility to correlate variables contained in different projects: for example the PS projects of the SCT and the cooling project of ID.
- possibility of dumping the entries in a text file.
- general tool that can be applied to any PVSS project without having to change the code.

None of the available tools were satisfying all these requirements, so we endeavoured in writing the AtlasDcsWebViewer.

To these initial requirements, more were added during development:

- possibility of overlapping trends from several data point elements
- possibility of histogramming the max, min and average value in a given time interval for a list of data points.
- possibility of navigating the Alias and data-point structure
- possibility of bit-AND (bit masking) operations on the results

- possibility of query on value bits (BITAND)
- wild-card searches for data points
- storing pre-defined lists of data points, or results of a search.
- possibility of pre-defined queries (like: all data-points related to a given ID cooling loop)

This note is primarily an introduction to a user manual, illustrating the various features of this web-available tool. There is also an Atlas Twiki page manual which is constantly updated. The second part of this note describes the structure of the project and its components.

2 User Manual

2.1 Supported browsers

At present the AtlasDcsWebViewer only supports Firefox. Other browsers have some intrinsic limitation, while for others there is a simple matter of manpower.

2.2 URL

The AtlasDcsWebViewer is available at the following URL

```
https://atlas-ide-dcs.web.cern.ch/atlas-ide-dcs/IdeDcsWebViewer/Interface/interface.html
```

The user is requested to be identified with his/her NICE account password.

The following **Atlas Twiki page** contains the up-to-date documentation:

```
https://twiki.cern.ch/twiki/bin/view/Atlas/IdeDcsWebViewer
```

2.3 Choosing the Detector Area

The AtlasDcsWebViewer can address the queries to all PVSS subsystems. However, for reasons of performances in navigation, we have subdivided the detector into four “areas of interest”: Inner Detector, Calorimeters, Muons and TDAQ. The basic idea is that it is rare that a user needs to correlate measurements across these systems. When this is needed one can afford to load in sequence the navigation menus of different parts. The

data points of the central DCS systems appear in all davigation menus. The LUCID subdetector has been arbitrarily assigned to the ID. By clicking on the desired part the user loads the navigation variables to find the data-point-element of interest: the navigation can be performed by *alias* or by *data-point-element name* (DPN). *Alias* and *DPN* navigation, as well as the different parts of the detector, can be mixed, with loading other parts in sequence. The loading of each navigation file (XML) should take no more than a few seconds.

2.4 Choosing the Time Interval

All searches on the DCS Oracle database must be limited within a time interval. Open-ended searches are not allowed. The start time and the end time are set with a calendar field. All database times are in GMT (CET-1). We have chosen to include a limitation on the number of values that can be retrieved for each element. Presently up to 15000 values can be retrieved, which imposes some limit to the maximum time interval that can be chosen without truncation of the results. The exact amount of time will depend on the frequency each element is written to the DB.

2.5 Selecting One Data Point Element

If the user knows already the data point element she wants to reach, she can do so by navigating through a selection of menus, which will appear in sequence on the right-hand side of the web page. If data-point navigation is selected the first pull-down menu to appear lets the user choose the subdetector, then a second menu appears to chose the PVSS project, then other menus will appear until the level of the data point element is reached. This way the user will have unambiguously specified one data point element. This element can be added to a list (which is an empty list if it is the first element), if the user wish to display more than one element at the time, or else the query can directly be sent to the database, to retrieve and plot the data for that particular data point only, as a function of time, within the specified time interval.

In the case of *Alias* navigation, the first submenu choice selects the subdetector, while the second submenu allows to select a group of projects. For instance, in case of the SCT the choice is between ENVironmental and PS (Power Supplies). The user navigates through the *Aliases* according to the hierarchy which is used in the subdetector.

2.6 The Display

To display the values as a function of time we use the freeware java-based tool “JFreeChart” [1]. Java needs to be installed as a plug-in in your Firefox browser. The plot will appear at the bottom of the web page only after some data are retrieved from the database.

Multiple variables can be displayed in the same plot, and there is a selection for using the same vertical scale or one scale per variable (element) in the display.

It is possible to use the mouse to zoom in and out of regions of interest. Right-clicking on the plot the user can save the plot, print it, zoom-in and out. The horizontal time scale is in UTC (GMT).

2.7 BitAND in Display

Some subdetectors use and store datapoints which encode status words. For instance, in the case of the SCT the state of a given channel is encoded in the *.State* datapoint: the high voltage state is coded in the four most significant bits, while the LV state is in the lower nibble. In this case it makes very little sense to plot as is the *.State* element. The high voltage state is selected with a bit-AND with 0x00F0, while the ON state is plotted with a bit-AND with The constant which is AND-ed bit by-bit with each value can be entered in decimal, or in hexadecimal or in binary.

2.8 Making a list of Data Point Elements

A list of *elements* will allow the user to display the values as a function of time for more than one element at the same time. Moreover, the user will be able to produce histograms of the average, the maximum values or the minimum values of the selected elements in the list. There are, at present, five ways to make a list of *elements* (or *aliases*):

- Selecting the elements or aliases one by one and clicking on “Add to the existing list”
- Using the free-space filter
- From special queries, like the “Loop query”.
- Loading a list from the “Personal Settings”.
- As a result of a “relational” query (query on values).

The first four methods allow the user to either create a new list or to add to the existing list of DPE’s, while the query on values always generates a new list. The list of elements can also be saved for later use, by using the commands in the “Personal Settings” tab.

2.9 Using the Free-Space Filter

The Free Space Filter tab allows the user to type in the full *element name* or *alias* or *comment* (description), or just a part of it and use “wild cards” to obtain a list of all the

elements, or the aliases or the descriptions that match the pattern. For instance, the user can simply select “*Description Filter*”, enter “Dew*” in the Free-Space filter and obtain the list of the 442 datapoints within all Atlas systems which measure the Dew point. The filter is case-sensitive, so “DEW” is different from “Dew”. It is also possible to transfer a DPE or alias name from the navigation menu or from the pending list of elements to the Free-Space Filter. This is useful for finding all the elements that are similar to a given one, without typing in the largest part of the name, with the risk of mis-typing.

It is possible to include within the same search more than one pattern by using the “or” operator || between patterns. It is also possible to exclude a given pattern from the result, by using the “veto” operator “-”. As an example the search on “IDE*loop001*power*—workingSensors” will retrieve all the data point elements that contain “power” but not “workingSensors”.

2.10 Keeping and sharing the lists

The list of elements can be kept for further usage. For instance a user can select a standard set of elements to monitor every day. She can store the pending list using the *Personal Settings* tab, by entering the name of the configuration and clicking on “Save Selected Items”. The configuration can be retrieved for later usage by selecting it among the list of the user’s configurations, which is obtained by clicking on “Get existing configurations”. Configurations are centrally stored at the server site. They can be shared among users. In order to do this it is sufficient to know the NICE account name of the other user, type it in the user field and click on “share”. This way the same list of elements will be copied in the new user’s “list of lists”, with the same apparent name as the original.

2.11 Queries on values

The queries on values are sometimes referred to as “Relational queries” and the title of the tab reflects this floppy slang. The queries of this type have as a result a list of datapoints whose values fulfill some requirements during the specified time interval. For instance, one can request the list of all SCT modules whose temperature had at least one value above 30 deg C within the time interval. This type of query is of the “out of range type”, while another possible query is, for instance, the list of all modules which had the HV reading always in the range 50 to 150 V. This is the “in-range” query. With the “Step” query we can select all the datapoints of a given type which had a step variation in their value. The step can be set to be above a minimum but below a maximum step. This is important if one wants to look for steps in the Low Voltages, which can indicate a malfunctioning, but at the same time avoid all the ON-OFF transitions, which are recorded as very large steps.

The query on values is somehow similar to the “freespace Filter”, which returns all the *DPE*s (or *aliases*) which match a given string pattern in their name. The query on values

adds to this some conditions on the values during the selected time interval. The starting point of a “query on values” is a uniquely identified data point element. The user then makes it general, requiring “all those like that one” that fulfill the request on the values.

From the practical point of view we can start with selecting one DPE (as described in section 2.5). We then select the “*Relational Query*” tab and click on “UPDATE relational query fields”. The element name (or alias) will be loaded to the area, and it will be subdivided in various fields.

The first field indicates the subdetector, and can not be subject of wild card search for matters of efficiency. This is probably a limitation, but in general the queries on values are limited to a given subsystem. The last field indicates the element, and again it can not be substituted with a wild card, because for instance one wants to limit the search to all the digital Low Voltages, i.e. only the elements of the same kind. These constraints are in place to avoid the possibility of launching a general query over the whole PVSS database on ANY datapoint. Everything that is in between the Subdetector specification and the element can be substituted with a wild card by pulling in the corresponding menu and selecting the “All” option.

Next step is to decide what kind of filter to apply. The possibilities are “in-range”, “out-of-range”, “Step” and “Bit-AND”. For the first 3 selections the user is required to type in two values, which are the lower and upper threshold values.

2.11.1 Out-of-range queries

The user has to first define the required subset of *data point elements* (or *aliases*) with the appropriate “wild card”, as described above. After selecting the Out-of-range query the user has to type the upper and lower threshold values. Clicking on “Make a relational query” sends the request to the Oracle Data Base. The query will return a list of all the datapoint elements (or *aliases*) which match the string and have at least one value above the upper threshold or below the lower threshold, within the specified time range.

2.11.2 In-range queries

This query will return a list of all the datapoint elements (or *aliases*) which match the string and during the specified time range had all the values inside the interval defined by the two thresholds.

2.11.3 Step queries

A step is defined as the difference between two consecutive values of the same datapoint element. This query will return a list of all the datapoint elements (or *aliases*) which match the string and during the specified time range had at least one step whose abso-

lute amplitude was within the range specified by the two thresholds. It only applies to elements which have at least 2 readings within the specified time range.

2.11.4 BitAND in queries on values

When selecting the Bit-AND option the user just needs to input one value, which is entered in decimal. This value will be compared bit by bit with the values in the archive for the set of elements (or aliases) which match the string. All the elements which have at least one value which gives a non-zero result of this bit-and operation will be reported in the result list. As usual, the test is performed on all the values which belong to the time interval of interest.

2.12 Special Queries

Special queries are detector-dependent and usually are just a short-hand for some standard “free-space” queries. Two “special queries” are presently implemented: one will return all the data point elements that are specific of a given Inner Detector cooling loop. The other will return all the elements of a given type for all the Inner Detector Cooling Loops (for instance all the elements which measure the heater surface temperatures). More queries of this kind can be added on request.

2.13 Histograms

The user may also want to obtain some statistical information about a system. For example, she may want to display the temperature distribution of a given part of subdetector for the specified time interval. In order to do so, it is necessary to have just one representative value per element, which can be the max, min or time average value. The starting point for the histogram is a list of elements or aliases, where we have selected just the elements of interest. On clicking on the “Generate Histogram” button a query will be sent to the Oracle DB. For each selected element of the list this query will return the maximum value, the minimum value and the time-weighted average for the elements which have at least one entry during the time interval. These “representative” values will also appear in the tooltip text of the *element* or *alias*. The user can choose to display the histogram of any combination of these 3 quantities. Then we can select all the elements, or just a subset of them, choose the number of bins in the histogram and click on “show histogram”. The histogram will appear in the lower part of the browser window. We can include or exclude elements within the list without querying again the database. The range is automatically selected by the display applet, while the number of bins can be set by the user.

2.14 Results

The results are available in a variety of forms. The results in alphanumeric format are available in the “Results” tab after performing a query that retrieves the time evolution of the *DP elements* (or *alias*).

For each element a sub-tab is present. Clicking on the sub-tab a list appears, with the time-ordered values in the format <value> <time-stamp>.

2.15 Saving the results (numeric)

The results can be saved by sending a file to your preferred e-mail addresses. It is possible to send the result to a comma separated list of e-mail addresses. The user can chose to send the file as attachment or as message body. At present the format can be txt, csv or xml, but we are planning to add a root ntuple format as well. All the available results are sent via e-mail. The first 3 lines (records) contain the *data point name*, the *alias* and the *description*, respectively. The second or third records can be empty if the *alias* or the *description* have not been set for that data point. All times are in UTC.

The following example contain the first few records for a data point which has no *alias* set.

```
ATLCICSDX1:DewPoint_SDX1_L2.value
```

```
CIC Environment SDX1Level2 Ambient Dew Point
3.134 01-12-2008 12:17:42:402
3.327 01-12-2008 12:18:42:403
```

The attachment file name has the following form:

```
<user>_IdeDcsWebViewer_Results_<Date>_<Time>.<type>
example:
```

```
dauria_IdeDcsWebViewer_Results_26_February_2009_10_56_26.txt
```

2.16 Saving the results (graphic)

The way of saving the graphic plot depends on the browser and operating system. In case of Linux, the suggested way is to use the display package and selectteh relevant window to copy the graph.

In case of Windows, right clicking on the plot will open a menu, where the option “print” can be selected, and “print to file” will produce a pdf file.

In case of Apple computers the user’s preferred option is teh screenshot.

2.17 Configuration and Personal Settings

The “Personal Settings” tab opens an area where the user can save, delete and share his/her lists of elements or aliases.

2.18 Troubleshooting

The “Error console” from the “Tools” tab of Firefox browser is usually a good tool to understand if there is something wrong. Typical hiccups are due to the java applet installation in the browser. The present version of the AtlasDcsWebViewer is 2.02. When loading a new version is always recommended to close completely the browser and clean the cache.

3 Architecture

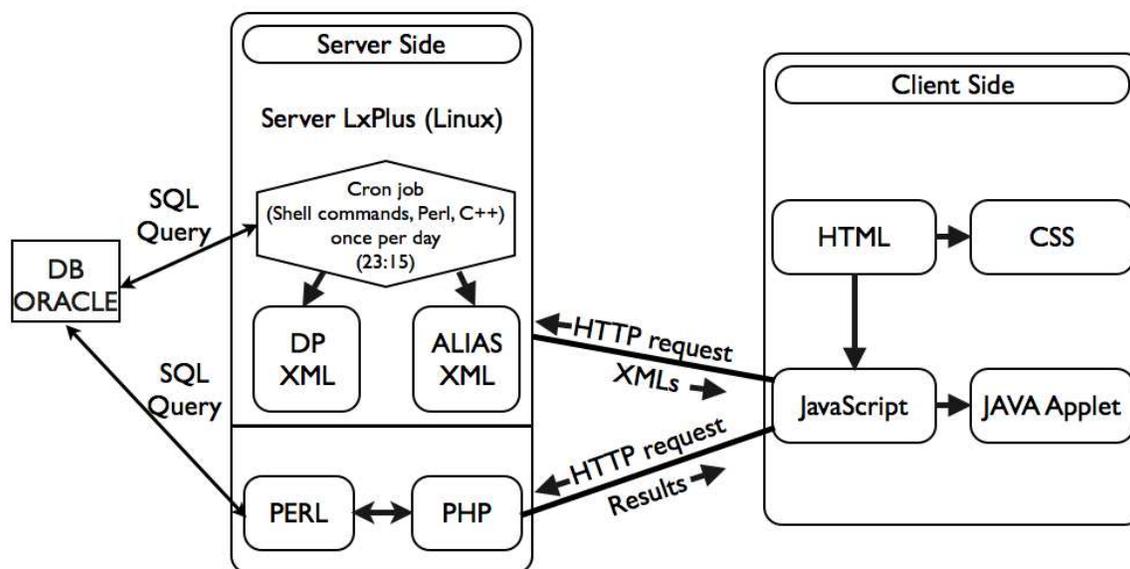


Figure 1: Architecture of the project

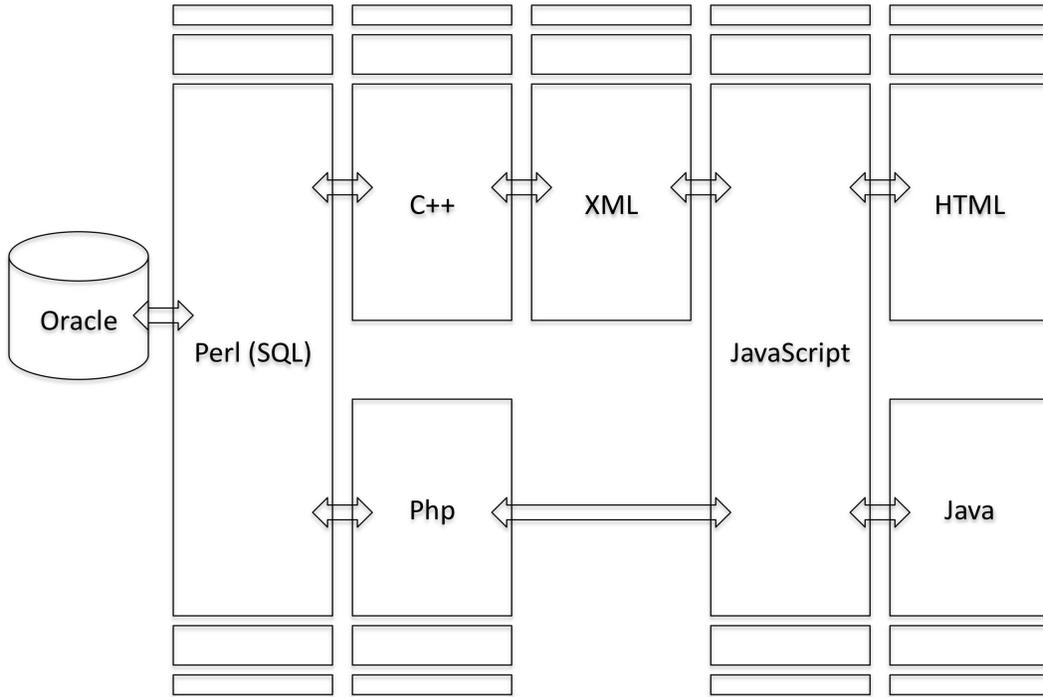


Figure 2: Summary scheme of the computing languages used in the project



Figure 3: Screenshot of the web page showing a list of elements

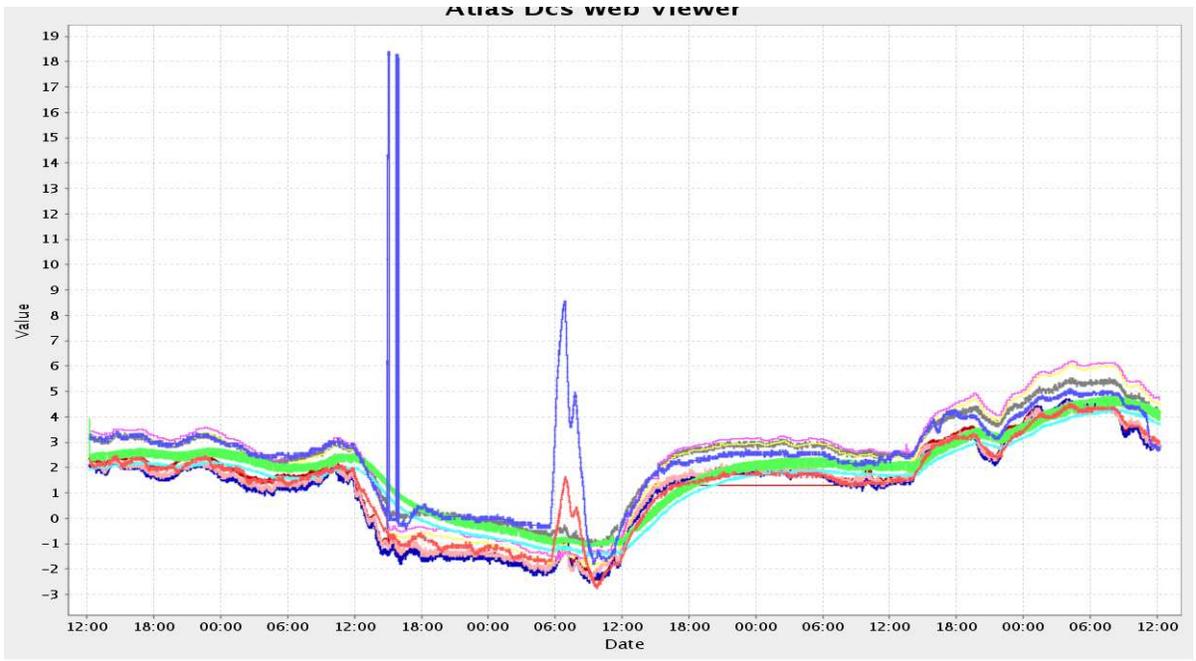


Figure 4: Example of a trending plot

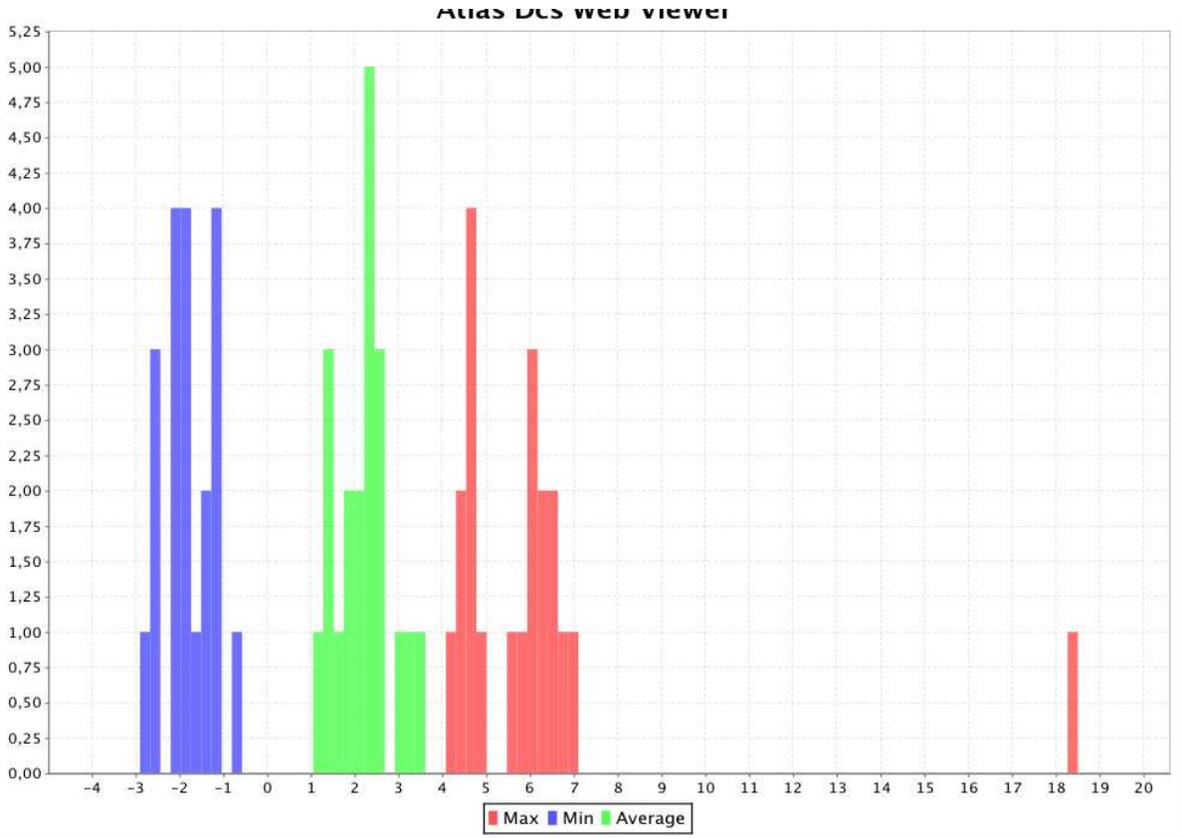


Figure 5: Example of a histogram, showing the distribution of the maximum, minimum and average values of the selected data point elements.

References

- [1] JFreeChart: <http://www.jfree.org/jfreechart/>
- [2] PVSS is produced by ETM GmbH, Kasernenstrae 29 A-7000, Eisenstadt, Austria,
<http://www.etm.at/>