

The *myStats* User's Guide

Version 2.1; January 2012

INTRODUCTION	1
COMMAND LINE OPTIONS	1
HELP TEXT	2
USER CONFIGURATION FILE	3
THREADED EXECUTION	3
PROGRESS REPORTING	3
OUTPUT	4
SAMPLE OUTPUT	4

Introduction

The *myStats* utility provides the means for users to compute statistics on EPICS channel history stored in the *Mya* archiver. The statistics include minimum, maximum, mean, and standard deviation. The tool runs on Redhat Enterprise 6 Linux systems that are connected to the Jefferson Lab accelerator division's network.

You specify a time span of interest and a set of EPICS channel names. In addition, you may specify one channel as the "control channel", which governs which periods within the user's time span are useful for statistical purposes. The useful periods are those where the control channel's value is within a range specified by the user. For example, the user may specify that statistics will only be computed for periods where beam current is between 100 and 500 micro-amps. Note that additional channels are optional and the tool may be used to gather statistics only for the control channel. When no control channel is specified, statistics on the other channels will be taken over the entire user entered time span.

Command line options

The *myStats* tool is run from a Linux shell with one or more command line options. You can get a summary of these options by entering '*myStats -h*' to the shell prompt. The following text box shows the tool's command line help. As can be seen from the help text, a user may specify program options from a configuration file or directly from the command line. An example of a user configuration file follows the help text box.

When options are given directly on the command line, channels of interest are specified using the '-l' option. The qualifier for this option may be a comma separated list, or the name of a file that contains the channel names; one per line. A file is designated by using the '@' as a prefix to the file name.

Help text

Version 2.1 - January 2012

To view the full User's Guide enter 'myStats -help'

There are two ways to provide arguments to this application; by configuration file or directly on the command line. The format of the begin date is 'YYYY-MM-DD HH:MM:SS', with the time portion optional.

Use quotes if date and time are entered due to the space separator.

Note the end date may be relative to the begin date when the '^<n><u>' format is used. The <n> is a number and <u> is a units specifier; (s,m,h,d,w) is for (seconds,minutes,hours,days,weeks). You may also enter an end date in the same format as the begin date. The format of the configuration file may be found in the User's Guide.

The output is written to standard output if an output file is not specified.

<< CFG >>

Specify by configuration file.

Usage: -cfg<opt> [-m<opt>] [-P] [-f<opt>] [-t<opt>]

cfg; User configuration file

m; Mya deployment; default = ops

P; Show progress

f; Output file

t; Number of threads; max = 10; default = 1

Example: myStats -cfg ./myInfo.cfg

<< Command-line >>

Specify directly on the command-line.

Usage: -b<opt> [-e<opt>] [-c<opt>] [-r<opt>] [-l<opt>] [-p<opt>] [-m<opt>] [-P] [-f<opt>] [-t<opt>]

b; Begin date

e; End date; default = ^1h

c; Control channel

r; Range of control (low:high)

l; Channel list (comma separated or @file)

p; Numeric precision for output

m; Mya deployment; default = ops

P; Show progress

f; Output file

t; Number of threads; max = 10; default = 1

Example: myStats -b '2009-09-14 10:35:00' -l R121PMES,R122PMES

Example: myStats -b -1h -e ^15m -c R2XXITOT -r 60:800 -l @myFile.txt

User configuration file

```
# Specify the date span of interest.
Begin = {2010-04-01 00:00:00}
End = ^1d # One day after begin

# Specify the channel that controls valid periods within the date span.
# Note these two keywords are optional.
ControlChannel = R2XXITOT
ControlRange = 60:800 # Envelope in uAmps.

# Optionally specify the numeric output precision. Remove or comment out the
# following line to have the default precision.
#Precision = 10

# Specify the channels that will be analyzed during valid periods. This is a
# comma separated list, which may be repeated as necessary. Note that no
# white space is allowed embedded in the list.
Channels = R121PMES,R122PMES,R123PMES,R124PMES
Channels = R125PMES,R126PMES,R127PMES,R128PMES
```

Threaded execution

The user may choose to perform data fetches and statistical calculations in parallel to obtain a performance boost. The application has been observed to run significantly faster using this option in certain execution scenarios. The command line offers a ‘-t’ option for this purpose. The default value of one means all processing will be done serially. Be aware, if you choose to process concurrently, the resulting output of channel data will no longer be in the order in which they were specified. Use just one thread when the output ordering is important to you. The suggested command line for long runs with many channels is as follows. Note you should try different thread counts to get a feel for how they affect your scenarios. There is no need to specify the maximum allowed when there is no performance increase with numbers greater than three or four.

myStats -cfg <your file> -P -f <your file> -t <n>

Progress reporting

Use the ‘-P’ option to enable progress reporting to standard output. This is useful for long running executions of the application. The output is a running count of the number of channels completed. It will always show you the execution time when complete. If you choose this option, it is best not to capture the calculated results by redirecting standard output to a file; because you will also capture all of the progress updates. Either let the entire output go to standard output or explicitly choose an output file using the ‘-f’ command line option.

Output

The *myStats* tool prints its output as a table to the specified file, or to “standard out. The output consists of one line per channel, having the statistics for that channel’s history over the time span. The first row is always the control channel, if one was provided. The table of data is nicely formatted for readability. The following output example was obtained by running the tool with the user configuration file shown above.

Sample output

Name	Min	Mean	Max	Sigma
R2XXITOT	60	73.2915	127.74	6.57701
R121PMES	149.1	151.341	155.8	2.63032
R122PMES	-92	-90.0156	-86.1	2.33463
R123PMES	-21.9	-19.9053	-15.8	2.41062
R124PMES	116.2	118.297	122.5	2.45755
R125PMES	52.1	53.9983	57.9	2.29101
R126PMES	-179.9	-177.927	-173.9	2.3502
R127PMES	-83	-80.9397	-76.9	2.41078
R128PMES	125.3	127.371	131.5	2.45878