



# LProf

Version 1.9  
April 2023

**MAQAO Tutorial series**

# 1 Introduction

MAQAO Lightweight Profiler (LProf) is the MAQAO module which allows to easily profile an application to detect hot functions and loops in two steps:

## 1) Data collection using sampling

LProf uses hardware counters to profile large-scale parallel applications (2000+ cores) with a very low overhead.

It is also possible to provide a custom list of hardware counters to sample.

## 2) Data display

It allows to **simply (avoid repeating “easily”?)** identify time-consuming functions and loops, observe the amount of time spent by the application between different categories (I/O, Runtime, etc...) and detect load balancing issues.

## 2 Running MAQAO LProf

### 2.1 Sequential Run Command

```
maqao lprof -- <application> [arg1 arg2 ...]
```

application's name (or path if not located in the current directory)

application's arguments, if any

### 2.2 Parallel Run Command

WARNING! Invocation of LProf in MPI has changed in old versions of MAQAO.

MAQAO Version 2.4 and below: the command line begins with the MPI command.

```
mpirun -n <NB_PROCESSES> maqao lprof -- <application> [arg1 arg2 ...]
```

MPI launcher command

number of processes

From 2.5 version and higher: the `--mpi-cmd` option is required.

Interactive runs:

```
maqao lprof --mpi-cmd="mpirun -n <NB_PROCESSES>" -- <application> [args]
```

MPI launcher command

number of processes

Batch runs:

```
maqao lprof --batch-script=<jobscript> [--mpi-command="mpirun..."] \
[--batch-command=<submission command>] -- <application> [args]
```

Required only if jobscript extension (e.g .sbatch) is not recognized

In jobscript, application executable and its arguments have to be replaced by `<run_command>`.

```
$ cat jobscript.sh
...
mpirun -n 4 <run_command> # instead of mpirun -n 4
<application> [args]
# <mpi_command> <run_command> # if mpi-command used
```

Since 2.12.0, you can (and must) inform Lprof about the maximum number of processes per node (if greater than 1), allowing it to set correct internal settings: `--maximum-processes-per-node`

Starting from 2.14.5, it is autodetected when missing but it is still recommended to set `--maximum-processes-per-node` if known and  $> 1$ .

## 2.3 Kernel samples exclusion

Since 2.12.0, kernel samples are not collected by default (recent Linux distributions do not allow this by default). To collect them:

- If `sysctl kernel.perf_event_paranoid` returns 2 or more, this step must be performed first:

```
$ sudo sysctl -w kernel.perf_event_paranoid=1
# lost after reboot
$ sudo sh -c 'echo kernel.perf_event_paranoid=1 >>
/etc/sysctl.d/local.conf
# persists after reboot'
```

- If `sysctl kernel.perf_event_paranoid` returns 1 or less:

```
$ maqao lprof --include-kernel ...
```

## 2.4 Options (collect step)

To list all options **along with their descriptions**:

```
maqao lprof --help
```

Options in gold color can be used to mitigate sampling overhead.

Options in light red can be used to override default behavior and workaround profiling issues.

Main options (collect step)		
Name	Short Description	Values
--include-kernel	No effect with the 'no-perf' engine. Count kernel samples (requires perf-event-paranoid level 1 or less)	(no value)
-mc/--mpi-command=	Specify command for interactive MPI run or replacement value for <mpi_command> in job script	Ex: "mpirun -n 4"
-bs/--batch-script=	Jobscrip to submit to job scheduler	Path to jobscrip (string)
-bc/--batch-command=	Command used to submit jobs, required if jobscrip extension is not recognized. Currently recognized: .sbatch and .pbs	Ex: "sbatch"
--stdin-path	Defines a file for redirection to stdin.	path to a stdin-redirection file
--sampling-rate=	Number of collected samples per second	<ul style="list-style-type: none"> <li>highest (2000 Hz, btm=off recommended)</li> </ul>

		<ul style="list-style-type: none"> <li>high (1000 Hz, avoid btm=stack)</li> <li><b>medium (200 Hz, default)</b></li> <li>low (50 Hz)</li> <li>lowest (10 Hz)</li> </ul>
<b>-ldi=</b>	Scan debug information into all or specified (provided list) library(ies) to get loops details	on (all)   <b>off (default)</b> or r list of libraries ('lib1, lib2, ...')
<b>-ug=</b>	Control ( <i>i.e.</i> pause/resume) measurement via a signal (Ctrl+Z) or via a countdown	on (CTRL+Z)   <b>off (default)</b> or a delay in seconds
<b>-btm=</b>	Select backtraces (callchains) collection method	<ul style="list-style-type: none"> <li><b>fp (default, recompile application with -fno-omit-frame-pointer)</b></li> <li>stack (higher overhead but no need to recompile application)</li> <li>branch (not really callchains but branch history, HW-dependent)</li> <li>off (no callchains, lowest overhead)</li> </ul>

### Advanced/other Options (collect step)

Name	Short Description	Values
<b>--use-OS-timers</b>	Use OS timers instead of hardware events. Needed in case of unavailable HW counters or undetected processor.	(no value)

	With autotuning features	
<code>--cpu-clock-MHz</code>	[perf-* engines] Override the "cpu-clock" perf-event rate (in MHz) measured by a calibration loop.	integer value
<code>--ref-cycles-MHz</code>	[perf-* engines] Override the "ref-cycles" perf-event rate (in MHz) measured by a calibration loop.	integer value
<code>--replace</code>	Overwrites an already existing output directory (reuse it). Remark: no effect on a not yet existing directory.	(no value)
<code>-tpp/--maximum-threads-per-process</code>	[perf-high-ppn only] Maximum number of concurrent threads per process. Default is OMP_NUM_THREADS. Used to set buffers and files size.	integer value
<code>--ppn/maximum-processes-per-node</code>	Since 2.12.0, mandatory when using <code>--mpi-command</code> Optional but recommended starting from 2.14.5 if <code>ppn &gt; 1</code>	Ex on single node: <code>lprof mpi-command="mpirun -n 32" ppn=32</code>
<code>--maximum-buffer-megabytes</code>	Allow to override Lprof memory footprint (default is 50 MB per CPU)	Maximum amount per node (Megabytes)
<code>--maximum-tmpfiles-megabytes</code>	Limit total temporary files size to X Megabytes per node. Default is 100 MB per CPU (HW thread).	Integer value
<code>-e/--evts</code>	Provide custom list of events to sample (CF <code>maqao --list-events</code> )	<code>evt1_name@sample_period, ...</code> or <code>evt1_code@sample_period, ...</code>

<code>-p/--evts-profiles</code>	Use ready-to-use lists of events. Not yet supporting more than one profile.	string
<code>--max-callchain-length</code>	Maximum callchain length (default: 20), useful to reduce btm=stack overhead.	Positive integer
<code>--stack-size</code>	Size (in bytes) of stack to dump on samples (default: 8192). Using a smaller size (typically 4096) reduces profiling overhead but may cut (or lose) callchains. Using a bigger size (typically 16384) increases profiling overhead but should guarantee minimal callchains loss.	Positive integer
<code>--mmap-pages</code>	Overrides autotuned number of mmap pages for ring buffer payload.	Positive integer
<code>--collect-calls-info</code>	Collects source file/line information for callchain nodes (calls). To display them, add <code>--use-calls-info=on</code> at display step.	on (default)/off
<code>--engine</code>	Use another perf-events based sampling engine	<ul style="list-style-type: none"> <li>• <code>perf-low-ppn</code> (selected by default when perf-events are available with max 4 processes per node)</li> <li>• <code>perf-high-ppn</code> (selected by default when perf-events are available with more than 4 processes per node)</li> <li>• <code>no-perf</code> (selected by default when perf-events are not available)</li> </ul>



<code>--include-sleep-time</code>	[no-perf only] Include sleep time (walltime).	(no value)
<code>--keep-external-threads</code>	[perf-high-ppn engine only] Profile threads with a different command line than the monitored application.	on/off (default)
<code>--keep-indirect-threads</code>	[perf-high-ppn engine only] Profile threads that are not direct children of the monitored application.	on (default)/off
<code>-cpu/--cpu-list</code>	Set CPU affinity for the target process. Ex: 0,2 to use CPU0 and CPU2.	comma-separated list of integers
<code>--ignore-signals</code>	[no-perf and perf-high-ppn engines] Prevents signals from being interpreted as termination signals. Allows to adapt no-perf and perf-high-ppn to various runtimes. Remark: for ignored signals also specified in set-exit-signals or set-abort-signals, evaluation order is set-abort-signals, set-exit-signals and then ignore-signals.	comma-separated list of integers
<code>--set-exit-signals</code>	[no-perf and perf-high-ppn engines] Interpret signals as normal application exit. Allows to adapt no-perf and perf-high-ppn engines to various runtimes. Remark: for exit signals also specified in ignore-signals or set-abort-signals, evaluation order is set-abort-signals, set-exit-signals and then ignore-signals.	comma-separated list of integers

<pre>--set-abort- signals</pre>	<p>[no-perf and perf-high-ppn engines] Interpret signals as abnormal application exit. Allows to adapt no-perf and perf-high-ppn engines to various runtimes. Remark: for abort signals also specified in ignore-signals or set-exit-signals, evaluation order is set-abort-signals, set-exit-signals and then ignore-signals</p>	
<pre>--legacy- maps</pre>	<p>[ADVANCED] Use only if unknown functions coverage is high for executable or libraries. Collect maps via legacy method (out of perf-events) after &lt;legacy-maps&gt; milliseconds and fallback to them in case of unresolved addresses.</p>	<p>Positive integer (number of milliseconds)</p>
<pre>--maximum- CPU-time- intervals</pre>	<p>[ADVANCED] [perf-low-ppn and perf-high-ppn engines] Maximum number of per-thread CPU-time intervals. Allows to trace when and where (CPU) threads was running, and display them by adding -verbose at display step.</p>	<p>Positive integer (number of intervals)</p>

## 2.5 Collect step hints

In case of multiple application processes (typically MPI ranks), use `collect-calls-info=off` to limit LProf memory footprint when dumping to disk source file/line for each call listed in callchains.

## 3 Display

The two common display modes are text (default) and HTML.

### 3.1 Concepts

LProf relates *code regions* contributions to *system-levels*. User must then specify which code regions he is interested in and at which system level/granularity.

#### 3.1.1 Code regions (hotspots)

From bigger to smaller:

- *Application*: set of *modules*
- *Module*: set of *functions*
- *Function*: set of *loops*
- *Loop*: set of *blocks*
- *Block*: basic block (compilation concept)

#### 3.1.2 System levels


From bigger to smaller:

- *Cluster*: set of *nodes* (machines)
- *Node*: set of (system) *processes*
- *Process*: set of (system) *threads*
- *Thread*

### 3.2 Text Output

#### 3.2.1 Functions Hotspots

To display summary view (at cluster level):



```
##### HOTSPOTS SUMMARY #####
# Function Name      Module      Source Info      Time Average (%) | Time Min(s) [TID] | Time Max(s) [TID] | Time Average(s) #
#####
# _INTERNAL_25_src_kmp  libtomp5.so 35.63           0.00 [67607]      37.53 [67668]      16.62 #
# MPIDI_CH3I_Progress  libmpi.so.12.0 33.49           2.11 [67617]     122.21 [67582]     15.62 #
# binvcrhs             binary      solve_subs.f:206 8.02           3.05 [67651]      4.18 [67581]       3.74 #
# matmul_sub          binary      solve_subs.f:56  3.52           1.30 [67672]      1.87 [67615]       1.64 #
# z_solve             binary      z_solve.f:4      2.74           1.08 [67689]      1.50 [67607]       1.28 #
# compute_rhs         binary      rhs.f:4          2.31           0.82 [67680]      1.30 [67615]       1.08 #
# y_solve             binary      y_solve.f:4      2.17           0.78 [67659]      1.30 [67615]       1.01 #
# x_solve             binary      x_solve.f:5      2.07           0.70 [67659]      1.15 [67675]       0.96 #
```

Figure 1 - LProf Output: Summary View (Functions)

To display view for a lower system level, use -dn (resp. dp, dt) for node (resp. process, thread). For instance, to display thread view:



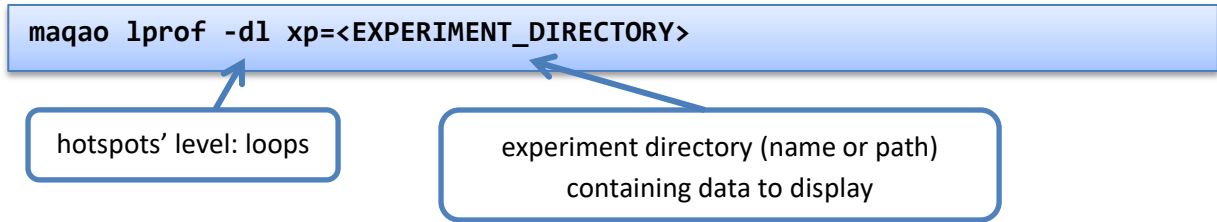
Thread ID      Hostname      Process ID      Walltime

```
##### CERES - PROCESS #67556 #####
##### Thread #67556 - 140.69 second(s) #####
# | Module | Source Info | Time % | Time(s) | CPI ratio | #
#####
# MPIDI_CH3I_Progress | libmpi.so.12.0 | 86.36 | 121.50 | 1.57 | #
# binvcrhs | binary | solve_subs.f:206 | 2.84 | 3.99 | 0.69 | #
# matmul_sub | binary | solve_subs.f:56 | 1.27 | 1.79 | 0.65 | #
# z_solve | binary | z_solve.f:4 | 0.97 | 1.36 | 0.73 | #
# compute_rhs | binary | rhs.f:4 | 0.87 | 1.22 | 0.80 | #
# x_solve | binary | x_solve.f:5 | 0.78 | 1.10 | 0.66 | #
# y_solve | binary | y_solve.f:4 | 0.73 | 1.03 | 0.60 | #
# matvec_sub | binary | solve_subs.f:27 | 0.40 | 0.56 | 0.89 | #
```

Figure 2 - LProf Output: Thread View (Functions)

### 3.2.2 Loops Hotspots

To display summary view (at cluster level):



```
#####
```

#	Loop ID	Module	Function Name	Source Info	Level
#	221	binary	matmul_sub	solve_subs.f:71-175	Single
#	230	binary	z_solve	z_solve.f:146-308	Innermost
#	227	binary	z_solve	z_solve.f:55-137	Innermost
#	204	binary	y_solve	y_solve.f:145-307	Innermost
#	197	binary	x_solve	x_solve.f:146-308	Innermost

```
#####
```

Figure 3 – LProf Output: Summary View (Loops)

The above figure is truncated. In the actual output, four more columns are available on the right (same as functions mode):

**Time Average (%)**, **Time Min (s)**, **Time Max (s)** and **Time Average (s)**.

As for functions, use -dn/dp/dt to select a lower system level. For instance, to display thread view:

```
maqao lprof -dl xp=<EXPERIMENT_DIRECTORY> -dt
```

```
#####
```

#	Loop ID	Module	Function Name	Source Info	Level	Time %	Time (s)	CPI ratio
#	221	binary	matmul_sub	solve_subs.f:71-175	Single	0.54	0.76	0.65
#	197	binary	x_solve	x_solve.f:146-308	Innermost	0.37	0.52	0.59
#	227	binary	z_solve	z_solve.f:55-137	Innermost	0.36	0.50	1.10
#	230	binary	z_solve	z_solve.f:146-308	Innermost	0.35	0.50	0.49
#	204	binary	y_solve	y_solve.f:145-307	Innermost	0.32	0.45	0.46
#	201	binary	y_solve	y_solve.f:55-137	Innermost	0.27	0.37	0.91
#	194	binary	x_solve	x_solve.f:57-139	Innermost	0.23	0.32	0.79
#	226	binary	z_solve	z_solve.f:415-423	Innermost	0.18	0.26	0.82
#	122	binary	compute_rhs	rhs.f:304-349	Innermost	0.16	0.23	1.01

```
#####
```

Figure 4 - LProf Output: Thread View (Loops)


### 3.3 Display Options

Basic Options (display step)		
Name	Short Description	Values
<b>-df/-dl</b>	Display functions/loops	(no value)
<b>-db</b>	Display basic blocks (for finer granularity than loops)	(no value)
<b>-dn</b>	Display per-node profiles (instead of cluster by default)	(no value)
<b>-dp</b>	Display per-process profiles (instead of cluster by default)	(no value)
<b>-dt</b>	Display per-thread profiles (instead of cluster by default)	(no value)
<b>-lec/--libraries-extra-categories</b>	Consider specified libraries as extra categories	libraries names as given by 'ldd <application>'
<b>-of/--output-format</b>	Output results in a file of the given format (default if omitted: console output)	html or csv
<b>-cc/--callchain</b>	Specify objects for callchains analysis: <ul style="list-style-type: none"> <li>• exe: display the callchain (if available) for each function with a scope limited to the application.</li> <li>• lib: extend the callchain scope to external libraries function calls.</li> <li>• all: display the callchain with no limited scope (application + libraries + system calls).</li> </ul>	exe, lib, all or off

	<ul style="list-style-type: none"><li>• off: disable callchains analysis. Some OpenMP/MPI functions/loops will no more be correctly categorized. Use this only when display takes too much time/memory.</li></ul>	
<b>-ct/--cumulative-threshold</b>	Display the top loops/functions up to a given cumulated coverage (e.g: ct=50).	integer between 0 and 100

## 3.4 HTML Output

### 3.4.1 Generation of HTML results



This command generates an 'index.html' file into the `<EXPERIMENT_PATH>/html/` directory. Open this file into a web browser to see the results.

### 3.4.2 Interpretation of the Results

Refer to the Oneview tutorial:

<https://maqao.org/documentation/MAQAO.Tutorial.ONEVIEW.pdf>