

1 The performance tool NestedEvents

A new profiling tool called `NestedEvents` was built as an extension of the existing tool `PetscLogEvent` which is part of the open source library PETSc. The new timer functionality will be sent to PETSc to be incorporated into new PETSc releases.

The new timers are activated by a call to `PetscLogInitializeNested` during the initialization, and calls to `PetscLogViewNested` and `PetscLogFreeNested` at the end of the program.

The software is accompanied by a test program, `SchoolDay.F90` (approximately 300 lines), which illustrates its usage and effect.

1.1 Example output

The nested timer functionality is demonstrated in Tables 1 and 2, which show the original and the new output for the sample program `SchoolDay`.

Table 1 provides a lot of statistics, but it does not show which event is part of other events. Hence, when all the times in the original timer output are added, the sum is much more than the total amount of time spent on the `SchoolDay`.

Table 2 shows contents of the file `performance.out`, which summarizes the performance measurements in three different ways:

1. Global information

Information about the entire run, identical to the information in the original output;

2. Nested information

Output in which certain events are shown as a sub-event of other events, for instance: `Correct Homework` is a sub-event of `Tidy Classroom` during the `Morning`. It is also clear that certain events can be part of several other events: `Lessons` happen during the `Morning` and during the `Afternoon`, and even outside these two time spans.

There is also information about the distribution of work among processes. All processes, for instance, model two `Mornings` (2 calls per parent call), but these mornings (combined) take 0.61 seconds for some processes, and 0.80 for others. During the `Morning`, some processes do and others do not `Play Ball` (0 or 1 calls per parent call).

The output is sorted per level, in order of consumed time.

`self-timer`

Note the events `self-Lessons`, `self-Morning`, `self-Afternoon` and `self-Tidy Classroom`. These events are presented for those activities for which sub-activities have been timed: it is the time spent inside the timed sections (`Lessons`, `Morning`, `Afternoon` or `Tidy Classroom`), but *not* in any of the timed sub-sections.

`other-timed`

A threshold is used for the time to determine whether an event should be printed or

ignored. If an event (e.g. `Lessons`) has multiple ignored sub-activities, these may add up to more than negligible amount. This amount is printed under the name `Lessons: other-timed`. The example-output has no items with `other-timed` parts.

3. Self-timers

The time spent in each event itself (and not in any of the timed sub-activities) is presented in the third 'view' of the output. Again, the timers are sorted in order of consumed time.

Though the output shows the diversity in what processes do, some information is lost about which process did what exactly. This information is available in separate output files named `performance_0000.out ... performance_0005.out`.

1.2 XML output

Apart from the plain text output to the log file, an XML file may be created with all performance statistics gathered in XML format. By default, the plain text format is selected, but this is switched to XML format by a call to the function `PetscViewerInitASCII_XML()` in between calls to `PetscViewerASCIIOpen()` and `PetscLogViewNested()`. Moreover, extra XML outputs can be written by calls to a series of `PetscViewerXML...` functions, as is shown in the demo program `SchoolDay.F90`. This alternative output is selected by a special initialization function, as can be observed in the code of `SchoolDay.F90`. A section of the XML output is shown in Table 3. Moreover, in case of a large nested tree, the plain text output is also not so human-friendly as one would like to collapse and expand the various sections of the tree.

One approach to convert an XML format to a more human-friendly form is by using an XSLT (Extensible Stylesheet Language Transformation) script. This is basically another XML file that contains rules to convert an XML file to another form, for instance to an html file that can be easily viewed in a web browser. Web browsers such as Firefox, Internet explorer and Chrome can automatically apply XSLT files to XML files. For this purpose, the script `performance_xml2html.xsl` (473 lines) has been developed that transforms the XML performance report to html. The html format also allows the functionality to collapse and expand the nested timer tree with the help of Javascript functions. The transform script works well the common internet browsers Firefox (tested on v38.2.0), Internet Explorer (tested on v11.0.22)) and Chrome (tested on v44.0). Chrome doesn't normally allow local transformation files to be processed. To circumvent this security setting, start Chrome with a command line option `--allow-file-access-from-file`.

If the xsl file is located in the same folder as the XML file, and if the files are located there as well, then the automatic conversion by the web browser results in a display such as Figure 1. The conversion to html can also explicitly be applied by the command line tool `xsltproc` (see <http://xmlsoft.org/XSLT/xsltproc.html>).

1.3 Acknowledgements

The code development that realize the features described in this manual has been funded by the Maritime Research Institute of the Netherlands, MARIN.

----- PETSc Performance Summary: -----
 SchoolDay.exe on a arch-linux2-c-debug named vortech35.vortech.nl with 2 processors, by bas Fri Sep 26 11:35:12 2014
 Using Petsc Release Version 3.5.2, Sep, 08, 2014

	Max	Max/Min	Avg	Total
Time (sec):	1.348e+00	1.00060	1.348e+00	
Objects:	1.000e+00	1.00000	1.000e+00	
Flops:	3.413e+06	3.27889	2.227e+06	4.453e+06
Flops/sec:	2.532e+06	3.27691	1.652e+06	3.304e+06
Memory:	3.550e+04	1.00000		7.101e+04
MPI Messages:	2.000e+00	1.00000	2.000e+00	4.000e+00
MPI Message Lengths:	4.800e+02	1.00000	2.400e+02	9.600e+02
MPI Reductions:	3.000e+00	1.00000		

Event	Count		Time (sec)		Flops		Mess	Avg len	Reduct	--- Global ---					--- Stage ---					Total Mflop/s	
	Max	Ratio	Max	Ratio	Max	Ratio				%T	%F	%M	%L	%R	%T	%F	%M	%L	%R		
ThreadCommRunKer	1	1.0	5.96e-06	1.0	0.00e+00	0.0	0.0e+00	0.0e+00	0.0e+00	0	0	0	0	0	0	0	0	0	0	0	0
ThreadCommBarrie	1	1.0	4.05e-06	1.3	0.00e+00	0.0	0.0e+00	0.0e+00	0.0e+00	0	0	0	0	0	0	0	0	0	0	0	0
MPI_iSend	2	1.0	3.60e-05	1.1	0.00e+00	0.0	2.0e+00	2.4e+02	0.0e+00	0	0	50	50	0	0	0	50	50	0	0	0
MPI_Recv	2	1.0	1.50e-01	5520.	0.00e+00	0.0	2.0e+00	2.4e+02	0.0e+00	6	0	50	50	0	6	0	50	50	0	0	0
MPI_AllReduce	2	1.0	2.00e-02	1336.	0.00e+00	0.0	0.0e+00	0.0e+00	2.0e+00	1	0	0	0	67	1	0	0	0	100	0	0
Morning	2	1.0	7.50e-01	1.1	2.10e+06	3.2	4.0e+00	2.4e+02	0.0e+00	52	62	100	100	0	52	62	100	100	0	0	4
Afternoon	1	1.0	3.70e-01	1.2	3.45e+05	1.4	0.0e+00	0.0e+00	2.0e+00	25	13	0	0	67	25	13	0	0	100	0	2
Play Ball	2	0.0	1.60e-01	0.0	0.00e+00	0.0	0.0e+00	0.0e+00	0.0e+00	6	0	0	0	0	6	0	0	0	0	0	0
Skip Rope	3	0.0	2.30e-01	0.0	0.00e+00	0.0	0.0e+00	0.0e+00	0.0e+00	9	0	0	0	0	9	0	0	0	0	0	0
Tidy Classroom	4	0.0	4.90e-01	0.0	2.14e+06	0.0	0.0e+00	0.0e+00	0.0e+00	18	48	0	0	0	18	48	0	0	0	0	4
Lessons	8	1.1	7.10e-01	1.2	6.74e+05	1.1	4.0e+00	2.4e+02	0.0e+00	49	29	100	100	0	49	29	100	100	0	0	2
Correct Homework	2	0.0	1.50e-01	0.0	4.69e+05	0.0	0.0e+00	0.0e+00	0.0e+00	6	11	0	0	0	6	11	0	0	0	0	3

----- End of PETSc Performance Summary -----

Table 1: Original PETSc timer output for test program SchoolDay. Some lines have been removed for compactness.

----- PETSc Performance Summary: -----

./SchoolDay.exe on a arch-linux2-c-debug named vortech37.vortech.nl with 6 processes, by refresco Mon Sep 14 12:25:46 2015
Using Petsc Release Version 3.4.3, Oct, 15, 2013

	Max	Max/Min	Average	Total
Time (sec):	1.360e+00	1.00061	1.359e+00	
Objects:	1.000e+00	1.00000	1.000e+00	
MFlop:	3.413e+00	3.27889	1.436e+00	8.617e+00
MFlop/sec:	2.510e+00	3.27893	1.056e+00	6.339e+00
Memory (MiB):	2.592e-02	1.00118	2.590e-02	1.554e-01
MPI Message Transfers:	2.000e+00	1.00000	2.000e+00	1.200e+01
MPI Message Volume (MiB):	4.578e-04	1.00000	4.578e-04	2.747e-03
MPI Reductions:	4.000e+00	1.00000		

Timer tree (time is in % of 1.36 s, threshold is 0.10 %):

Morning	44 - 58 %	(2.0 calls/parent call)	{0.85 Mflop/sec}	[0.00 MBytes/sec]
Lessons	25 - 37 %	(2.0 calls/parent call)	{0.57 Mflop/sec}	[0.00 MBytes/sec]
self-Lessons	25 %	(2.0 calls/parent call)	{0.84 Mflop/sec}	
MPI_Recv	0 - 12 %	(0.5 calls/parent call)	[0.00 MBytes/sec]	
Tidy Classroom	0 - 16 %	(0.0 - 1.0 calls/parent call)	{6.51 Mflop/sec}	
self-Tidy Classroom	0 - 13 %	(0.0 - 1.0 calls/parent call)	{6.66 Mflop/sec}	
Correct Homework	0 - 3 %	(0.0 - 0.5 calls/parent call)	{5.84 Mflop/sec}	
Play Ball	0 - 13 %	(0.0 - 1.0 calls/parent call)		
Skip Rope	0 - 12 %	(0.0 - 1.0 calls/parent call)		
self-Morning	7 %	(2.0 calls/parent call)	{3.78 Mflop/sec}	
Afternoon	22 - 28 %	{0.92 Mflop/sec}	<5.32 Reductions/sec>	
Tidy Classroom	0 - 12 %	(0.0 - 1.0 calls/parent call)	{0.56 Mflop/sec}	
self-Tidy Classroom	0 - 8 %	(0.0 - 1.0 calls/parent call)	{0.16 Mflop/sec}	
Lessons	0 - 4 %	(0.0 - 1.0 calls/parent call)	{1.44 Mflop/sec}	
Lessons	11 %	(2.0 calls/parent call)	{1.30 Mflop/sec}	
Play Ball	0 - 6 %	(0.0 - 1.0 calls/parent call)		
Skip Rope	0 - 5 %	(0.0 - 1.0 calls/parent call)		
self-Afternoon	4 %	{0.98 Mflop/sec}		
MPI_AllReduce	0 - 2 %	(2.0 calls/parent call)	<82.07 Reductions/sec>	
Tidy Classroom	0 - 8 %	(0.0 - 1.0 calls/parent call)	{5.56 Mflop/sec}	
Correct Homework	0 - 8 %	(0.0 - 1.0 calls/parent call)	{2.13 Mflop/sec}	
Play Ball	0 - 6 %	(0.0 - 1.0 calls/parent call)		
Lessons	5 %	{1.71 Mflop/sec}		
Skip Rope	0 - 5 %	(0.0 - 1.0 calls/parent call)		

Self-timers (time is in % of 1.36 s):

Lessons	41 - 45 %	{1.09 Mflop/sec}
Tidy Classroom	0 - 29 %	{4.57 Mflop/sec}
Play Ball	0 - 19 %	
Skip Rope	0 - 17 %	
MPI_Recv	0 - 12 %	[0.00 MBytes/sec]
Correct Homework	0 - 11 %	{3.12 Mflop/sec}
Morning	7 %	{3.78 Mflop/sec}
Afternoon	4 %	{0.98 Mflop/sec}
MPI_AllReduce	0 - 2 %	<82.07 Reductions/sec>

----- End of PETSc Performance Summary -----

Table 2: *Nested timer output for test program SchoolDay*

```

<timertree desc="Timings tree">
  <totaltime>1.42104</totaltime>
  <timethreshold>0.100000</timethreshold>
  <event>
    <name>Morning</name>
    <time>
      <minvalue>42.951172</minvalue>
      <maxvalue>55.723322</maxvalue>
    </time>
    <ncalls>
      <value>2.000000</value>
    </ncalls>
    <mflops>
      <value>1.0852427</value>
    </mflops>
    <mbps>
      <value>0.001172</value>
    </mbps>
    <events>
      <event>
        <name>Lessons</name>
        ...
      </event>
      ...
    </events>
  </event>
  ...
</timertree>

```

Table 3: Section of the nested timer output in XML format for test program SchoolDay.

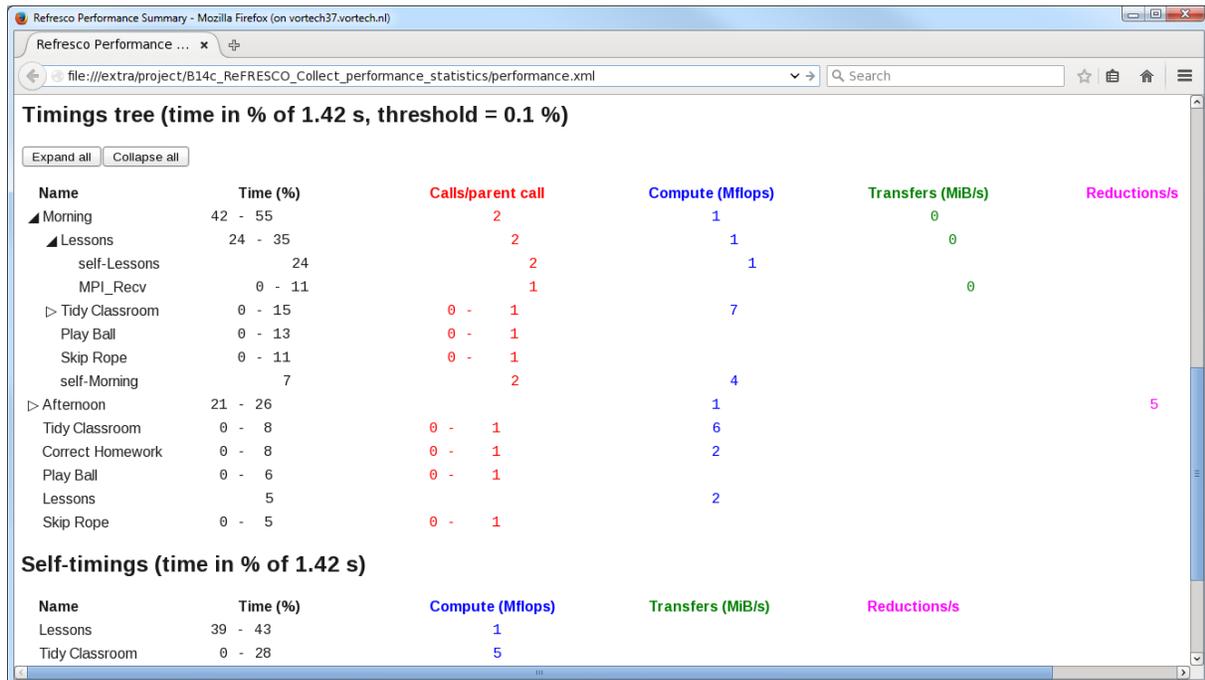


Figure 1: Transformed performance report in Firefox.