

# Package ‘pbdPROF’

July 2, 2014

**Version** 0.2-3

**Date** 2014-05-17

**Title** Programming with Big Data --- MPI Profiling Tools

**Depends** R (>= 3.0.0), methods, graphics, grid, ggplot2, gridExtra, reshape2

**LazyLoad** yes

**LazyData** yes

**Copyright** See pbdPROF/inst/fpmi\_LICENSE.txt for the files in src/fpmi/ and src/fpmi\_win.

**Description** MPI profiling tools.

**SystemRequirements** OpenMPI (>= 1.5.4) on Solaris, Linux, Mac, and FreeBSD. No MPI library required on Windows yet.

**License** Mozilla Public License 2.0

**URL** <http://r-pbd.org/>

**BugReports** <http://group.r-pbd.org/>

**MailingList** Please send questions and comments regarding pbdR to RBigData@gmail.com

**Author** Wei-Chen Chen [aut, cre], Drew Schmidt [aut], Gaurav Sehrawat [aut], Pragneshkumar Patel [aut], George Ostrouchov [aut]

**Maintainer** Wei-Chen Chen <wccsnow@gmail.com>

**NeedsCompilation** yes

**Repository** CRAN

**Date/Publication** 2014-05-26 22:39:57

## R topics documented:

pbdPROF-package	2
benchplot	3
Example datasets	4
Plot	4
Print	5
prof-class	6
Profilers	6
read prof	7
<b>Index</b>	<b>9</b>

---

pbdPROF-package      *MPI Profiling Tools*

---

### Description

This package contains several libraries for profiling MPI codes, as well as some R tools for parsing, analyzing, and plotting their outputs.

### Details

Package:    pbdPROF  
 Type:      Package  
 License:    MPL  
 LazyLoad:  yes

This package requires an MPI library (OpenMPI, MPICH2, or LAM/MPI).

### Author(s)

Wei-Chen Chen, Drew Schmidt, Gaurav Sehrawat, Pragneshkumar Patel, and George Ostrouchov.

### References

Programming with Big Data in R Website: <http://r-pbd.org/>

### Examples

```
## Not run:
demo(allreduce.fmpi, "pbdPROF")
demo(svd.fmpi, "pbdPROF")
demo(masterslavePI.fmpi, "pbdPROF")
demo(allreduce.mpip, "pbdPROF")
demo(svd.mpip, "pbdPROF")
demo(masterslavePI.mpip, "pbdPROF")
```

```
## End(Not run)
```

---

benchplot	<i>Benchmark Plotter</i>
-----------	--------------------------

---

## Description

Simple benchmarks plotter.

## Usage

```
benchplot(cores, timings, plot.type="speedup", title=NULL)
```

## Arguments

cores	Numeric vector containing the cores used in each benchmark.
timings	A dataframe containing the wall-clock timings of each benchmark.
plot.type	character; determines what kind of plot to produce. options are "runtime", "speedup", or "both".
title	character or NULL; plot title.

## Details

benchplot() produces simple benchmark plots. If multiple benchmarks are given (i.e., multiple columns are present in the timings dataframe) then the generated plot will automatically plot them against each other, separated by color.

## See Also

[prof-class](#)

## Examples

```
## Not run:
library(pbdPROF)

x <- c(1, 2, 4, 6, 8, 10, 12, 16)
y <- c(111.424, 80.696, 42.832, 29.468, 24.060, 19.390, 15.938, 11.860)

benchplot(x, y, "both")

## End(Not run)
```

---

Example datasets	<i>Small example datasets</i>
------------------	-------------------------------

---

### Description

Sample profiler outputs for testing and package demonstration.

### Format

Each dataset contains information profiled by **fpmpi** for pbdMPI/demo/allreduce.fpmpi.r, pbdDMAT/demo/svd.fpmpi.r, and Rmpi/demo/masterslavePI.fpmpi.r.

Each dataset contains information profiled by **mpip** for pbdMPI/demo/allreduce.mpip.r, pbdDMAT/demo/svd.mpip.r, and Rmpi/demo/masterslavePI.mpip.r.

### Details

Several sample profiler outputs are provided with the package, and utilized in the package demos to show off the functionality of the parsing and plotting methods provided by the package. These small datasets are located in the pbdPROF/extdata/ folder of the installed package (and pbdPROF/inst/extdata/ in the package source).

The example datasets are profiled in 2 processors.

---

Plot	<i>Plotting</i>
------	-----------------

---

### Description

Plot methods for prof class objects.

### Usage

```
## S4 method for signature 'prof'
plot(x, ..., which=1L:4L, show.title=TRUE,
      plot.type="timing", label, bar.label=FALSE)
## S3 method for class 'prof'
autoplot(object, ...)
```

### Arguments

x	prof class object
object	prof class object
...	extra arguments
which	Vector consisting of a subset of the integers 1, 2, 3, 4. Determines which plots will be produced.

<code>show.title</code>	Logical; determines whether or not the plot title will be shown.
<code>plot.type</code>	for profiling mpiP including "timing", "stats1", "stats2", "messages1", and "messages2".
<code>label</code>	The label for the plot title. Should be a character string or left missing for the default.
<code>bar.label</code>	logical; indicates whether or not numeric values of heights of barplots should be included should .

**Methods**

```
signature(x = "prof")
```

**See Also**

[prof-class](#), [read.prof](#)

---

 Print

*Printing*


---

**Description**

Print and show methods for prof class objects.

**Usage**

```
## S4 method for signature 'prof'
print(x, ...)
## S4 method for signature 'prof'
show(object)
```

**Arguments**

<code>x, object</code>	prof class object
<code>...</code>	extra arguments

**Methods**

```
signature(x = "prof")
```

**See Also**

[prof-class](#), [read.prof](#)

---

 prof-class

*Class prof*


---

### Description

Class for Profiler Output

### Slots

**profiler:** the profiler used stored as character data.

**raw:** raw profiler output (read in via `readLines()` cast as the virtual class `rawprof`).

**parsed:** a dataframe containing the parsed version of the raw slot.

### Details

The `prof` class is a simple container for managing profiler output via R.

The slots are `profiler`, `raw`, and `parsed`. The first, `profiler`, is just a character string consisting of the profiler used to generate the output (e.g., "fpmpi", "mpiP", etc.).

Next, the slot `raw` contains the raw output from the profiling library, stored in R as a vector of character data (strings) directly from a `readLines()` call, and cast as the virtual class `rawprof`. In reality, we will cast the raw data as a virtual class of the same name as that found in the `profiler` slot, and `rawprof` is a superclass of each of these classes. This is for internal S3 dispatch. You should not mess around with this.

The final slot, `parsed`, contains a condensed version of the information from the `raw` slot, stored as a dataframe. This is the representation used for plotting.

Unless you really think you know what you are doing, you probably shouldn't directly mess around with the elements of this class. Instead, simply use the `read.prof()` function.

### See Also

[read.prof](#)

---

 Profilers

*Profilers*


---

### Description

Profilers Directly Supported by pbdPROF

### Details

Currently, the **fpmpi** library is fully supported; a version of this library is bundled with the source of the pbdPROF package (see package vignette for more details). Additional libraries can easily be linked with pbdPROF, but these are not yet fully supported. The **mpiP** and **TAU** profilers are expected to be fully supported by the conclusion of Google Summer of Code (~September 30 2013).

## References

Programming with Big Data in R Website: <http://r-pbd.org/>

**fpmpi** website: <http://www.mcs.anl.gov/research/projects/fpmi/WWW/>

**mpiP** website: <http://mpip.sourceforge.net/>

**TAU** website: <http://www.cs.uoregon.edu/research/tau/home.php>

---

read prof

*Reading Profiling Outputs*

---

## Description

Reader for profiler outputs.

## Usage

```
read.prof(file.name, ...)
```

## Arguments

file.name	a full file name.
...	options for readLines

## Details

This function reads in profiling outputs from MPI-using R code and stores the output in a prof class object. The reading is managed by the `base::readLines()` function. The user does not need to specify the type of profiler output being used (e.g., whether the profiler text is from **fpmpi**, **mpiP**, etc.).

Additionally, this method automatically parses the output into a condensed, manageable dataframe (the `parsed` slot of the prof class).

## Value

A prof class object.

## See Also

[prof-class](#)

**Examples**

```
## Not run:  
library(pbdPROF)  
  
fn <- system.file("data/allreduce.fpmapi", package = "pbdPROF")  
da <- read.prof(fn, lib.type = "fpmapi")  
  
da  
  
## End(Not run)
```



# Index

- \*Topic **Classes**
  - prof-class, 6
- \*Topic **Methods**
  - Plot, 4
  - Print, 5
- \*Topic **Package**
  - pbdPROF-package, 2
- \*Topic **data**
  - Example datasets, 4
- \*Topic **utility**
  - benchplot, 3
  - read prof, 7

allreduce.fpmpi (Example datasets), 4  
allreduce.mpip (Example datasets), 4  
autoplot.prof (Plot), 4

benchplot, 3

Example datasets, 4

fpmpi-class (prof-class), 6  
fpmpi\_example (Example datasets), 4

masterslavePI.fpmpi (Example datasets), 4  
masterslavePI.mpip (Example datasets), 4  
mpip-class (prof-class), 6  
mpip\_example (Example datasets), 4

pbdPROF (pbdPROF-package), 2  
pbdPROF-package, 2  
Plot, 4  
plot (Plot), 4  
plot,prof-method (Plot), 4  
plot-method (Plot), 4  
Print, 5  
print (Print), 5  
print,prof-method (Print), 5  
print-method (Print), 5  
prof (prof-class), 6

prof-class, 6  
Profilers, 6

rawprof-class (prof-class), 6  
read prof, 7  
read.prof, 5, 6  
read.prof (read prof), 7

show (Print), 5  
show,prof-method (Print), 5  
show-method (Print), 5  
svd.fpmpi (Example datasets), 4  
svd.mpip (Example datasets), 4

tau-class (prof-class), 6