Package 'hpcwld'

February 20, 2015

Version 0.5

Date 2015-02-14

Title High Performance Cluster Models Based on Kiefer-Wolfowitz Recursion

Depends R (>= 2.10), multicool, partitions

Description Probabilistic models describing the behavior of workload and queue on a High Performance Cluster and computing GRID under FIFO service discipline basing on modified Kiefer-Wolfowitz recursion. Also sample data for inter-arrival times, service times, number of cores per task and waiting times of HPC of Karelian Research Centre are included, measurements took place from 06/03/2009 to 02/30/2011. Functions provided to import/export workload traces in Standard Workload Format (swf). Stability condition of the model may be verified either exactly, or approximately.

License GPL (>= 2)

URL http://www.r-project.org, http://cluster.krc.karelia.ru

Author Alexander Rumyantsev [aut, cre]

Maintainer Alexander Rumyantsev <ar0@sampo.ru>

Repository CRAN

Repository/R-Forge/Project hpcwld

Repository/R-Forge/Revision 11

Repository/R-Forge/DateTimeStamp 2015-02-17 14:13:21

Date/Publication 2015-02-17 19:32:19

NeedsCompilation no

R topics documented:

hpcwld-package 2
ApproxC 3
DataToSWF 4
DMC 5
FromSWF 6

HPC_KRC 7
HPC_KRC2 8
StabC 9
ToSWF 10
Wld 11
X 12

Index **13**

`hpcwld-package` *Model and data for High Performance Cluster workload*

Description

This package contains several models describing the behavior of workload and queue on a High Performance Cluster and computing GRID under FIFO service discipline basing on modified Kiefer-Wolfowitz recursion. Also sample data for inter-arrival times, service times, number of cores per task and waiting times of HPC of Karelian Research Centre are included, measurements took place from 06/03/2009 to 02/30/2011. The stability condition of the model can be verified either exactly, or approximately.

Details

Package:	hpcwld
Type:	Package
Version:	0.5
Date:	2015-02-14
License:	GNU GPL
LazyLoad:	yes

Author(s)

Alexander Rumyantsev (Institute of Applied Mathematical Research, Karelian Research Centre, RAS)

References

E.V. Morozov, A.Rumyantsev. Stability analysis of a multiprocessor model describing a high performance cluster. XXIX International Seminar on Stability Problems for Stochastic Models and V International Workshop "Applied Problems in Theory of Probabilities and Mathematical Statistics related to modeling of information systems". Book of Abstracts. 2011. Pp. 82–83.

A. Rumyantsev. Simulating Supercomputer Workload with hpcwld package for R // Proceedings of 2014 15th International Conference on Parallel and Distributed Computing, Applications and Technologies. IEEE, 2014. P. 138-143. URL: http://conferences.computer.org/pdcat/2014/papers/8334a138.pdf

A. Rumyantsev. Evaluating the stability of supercomputer workload model // Journal on Selected Topics in Nano Electronics and Computing, Vol. 2, No. 2, December 2014. P. 36-39.

http://cluster.krc.karelia.ru

Examples

```
Wld(T=rexp(1000,1), S=rexp(1000,1), round(runif(1000,1,10)), 10)
# returns the workload, delay and total cpus used
# for a cluster with 10 CPUs and random exponential times
```

`ApproxC` *Approximate, dynamic iterative computation of the stability constant for a workload of a High Performance Cluster model*

Description

This function calculates the constant C that is used in the stability relation of a model, which is basically the following: lambda/mu<C, where lambda is the intensity of task arrivals, and mu is the intensity of service, in this case it is 1/(mean service time). The constant depends only on the number of cores in the model and the distribution of core requirement of tasks. Note that this method of calculation allows to stop on some depth of dynamics, thus allowing to calculate an approximate value in faster time. The constant is valid only for the model with simultaneous task service by cores.

Usage

```
ApproxC(s, p, depth=3)
```

Arguments

`s`	number of cores/servers for a HPC
`p`	the distribution of number of cores required by a task, a vector with values for probability of a task requiring 1..s cores
`depth`	the depth of dynamical calculation. By default calculates up to groups of 3 tasks occupying the CPUs. When depth=s, calculates the exact value. However, depth=s might take a bit more time.

Value

The value of a constant C in the relation lambda/mu < C is returned

Author(s)

Alexander Rumyantsev (Institute of Applied Mathematical Research, Karelian Research Centre, RAS)

References

E.V. Morozov, A.Rumyantsev. Stability analysis of a multiprocessor model describing a high performance cluster. XXIX International Seminar on Stability Problems for Stochastic Models and V International Workshop "Applied Problems in Theory of Probabilities and Mathematical Statistics related to modeling of information systems". Book of Abstracts. 2011. Pp. 82–83.

Examples

```
StabC(s=2,p=c(.5,.5))
# returns the constant for a 2-server model,
# where a task needs 1 or 2 cores with equal probability
```

`DataToSWF` *Convertor from a dataframe to Standart Workload Format*

Description

Note that this is only a wrapper for the ToSWF command with a dataframe argument. It needs a correctly built dataframe and converts it to a Standart Workload Format used to share the logfiles of High Performance Clusters

Usage

```
DataToSWF(Frame, filename="output.swf")
```

Arguments

`Frame`	A dataframe containing the variables needed by ToSWF function.
`filename`	The file to store the converted workload (output.swf by default)

Details

The Standart Workload Format is a single format to store and exchange high performance cluster logs, that is used in Parallel Workload Archive. See references for current standard. The SWF format may contain additional data, but in this package only the 1st to 5th fields are used. One may also need to manually fill in the header of the file in order to completely prepare the resulting SWF file.

Value

Nothing is returned, but a file is created in the current working directory (with default name output.swf) containing the converted data.

Author(s)

Alexander Rumyantsev (Institute of Applied Mathematical Research, Karelian Research Centre, RAS)

References

Feitelson, D.G. and Tsafrir, D. and Krakov D. 2012 Experience with the Parallel Workloads Archive. Technical Report 2012-6, School of Computer Science and Engineering, the Hebrew University April, 2012, Jerusalem, Israel

http://www.cs.huji.ac.il/labs/parallel/workload/swf.html

See Also

FromSWF,ToSWF

Examples

```
##---- Should be DIRECTLY executable !! ----
##-- ==>  Define data, use random,
##--or do  help(data=index)  for the standard data sets.
data(HPC_KRC)
tmp=data.frame(T=HPC_KRC$interarrival, S=HPC_KRC$service, N=HPC_KRC$cores_used, D=HPC_KRC$delay)
DataToSWF(tmp)
```

DMC *Distributional Measure of Correlation*

Description

This is a suggested by Dror Feitelson measure of correlation for dependent variables, that may be successfully used to examine the datasets from a High Performance Cluster logs

Usage

```
DMC(X, Y)
```

Arguments

X	First variable (vector)
Y	Second variable (vector)

Value

One value between -1 and 1, characterizing the dependence between the variables

Author(s)

Alexander Rumyantsev (Institute of Applied Mathematical Research, Karelian Research Centre, RAS)

References

http://interstat.statjournals.net/YEAR/2004/abstracts/0412001.php?Name=412001

Examples

```
data(HPC_KRC)
DMC(HPC_KRC$service[1:1000], HPC_KRC$cores_requested[1:1000])
```

FromSWF *Convertor to a dataset from a Standart Workload Format*

Description

This is a convertor from a Standart Workload Format (used to share the logfiles of High Performance Clusters) to an internally used in a package dataset format

Usage

```
FromSWF(filename)
```

Arguments

filename	A mandatory field containing the path to SWF file

Details

The Standart Workload Format is a single format to store and exchange high performance cluster logs, that is used in Parallel Workload Archive. See references for current standard. The SWF format may contain additional data, but in this package only the 1st to 5th fields are used. One may also need to manually fill in the header of the file in order to completely prepare the resulting SWF file.

Value

A dataset is returned, containing 'delay' as a vector of delays exhibited by each task, 'total_cores' as the total busy CPUs in time of arrival of each task, and 'workload' as total work left at each CPU.

Author(s)

Alexander Rumyantsev (Institute of Applied Mathematical Research, Karelian Research Centre, RAS)

References

Feitelson, D.G. and Tsafrir, D. and Krakov D. 2012 Experience with the Parallel Workloads Archive. Technical Report 2012-6, School of Computer Science and Engineering, the Hebrew University April, 2012, Jerusalem, Israel

http://www.cs.huji.ac.il/labs/parallel/workload/swf.html

See Also

ToSWF,DataToSWF

HPC_KRC *Workload data for High Performance Cluster of High Performance Data Center of Karelian Research Center, Russian Academy of Sciences.*

Description

This is a complete data of the tasks which successfully finished executions at HPC of HPDC KRC RAS for time period 06/03/2009 to 02/04/2011, a total of 8282 tasks. The data contains interarrival times, service times, cores that tasks requested, cores really used (due to administrative limitations) and delays excursed by tasks, all in seconds.

Usage

```
data(HPC_KRC)
```

Format

A data frame with 8281 observations on the following 5 variables.

`interarrival` a numeric vector

`service` a numeric vector

`cores_requested` a numeric vector

`cores_used` a numeric vector

`delays` a numeric vector

Source

http://cluster.krc.karelia.ru

References

http://cluster.krc.karelia.ru

Examples

```
data(HPC_KRC)
```

HPC_KRC2 *Workload data for High Performance Cluster of High Performance Data Center of Karelian Research Center, Russian Academy of Sciences.*

Description

This is a complete data of the tasks which successfully finished executions at HPC of HPDC KRC RAS for time period 02/04/2011 to 16/04/2012, a total of 9389 tasks. The data contains interarrival times, service times, cores that tasks used, and delays excursed by tasks, all in seconds.

Usage

```
data(HPC_KRC2)
```

Format

A data frame with 9389 observations on the following 3 variables.

`interarrival` a numeric vector

`service` a numeric vector

`cores_used` a numeric vector

`delays` a numeric vector

Source

http://cluster.krc.karelia.ru

References

http://cluster.krc.karelia.ru

Examples

```
data(HPC_KRC2)
```

StabC *Stability constant for a workload of a High Performance Cluster model*

Description

This function calculates the constant C that is used in the stability relation of a model, which is basically the following: lambda/mu<C, where lambda is the intensity of task arrivals, and mu is the intensity of service, in this case it is 1/(mean service time). The constant depends only on the number of cores in the model and the distribution of core requirement of tasks. Note that there are three methods of calculation: two of them calculate the exact value ("less-memory" and "more-memory"), and the third one ("monte-carlo") returns the approximate value. The exact value may be calculated up to a number of cores of the order of 100. The constant is valid only for the model with simultaneous task service by cores.

Usage

```
StabC(s, p, maxiter=10000, method = "monte-carlo")
```

Arguments

s	number of cores/servers for a HPC
p	the distribution of number of cores required by a task, a vector with values for probability of a task requiring 1..s cores
maxiter	number of iterations for Monte-Carlo method
method	monte-carlo, less-memory, more-memory (approximately 1.5 times faster, but needs more memory)

Value

The value of a constant C in the relation lambda/mu < C is returned

Author(s)

Alexander Rumyantsev (Institute of Applied Mathematical Research, Karelian Research Centre, RAS)

References

E.V. Morozov, A.Rumyantsev. Stability analysis of a multiprocessor model describing a high performance cluster. XXIX International Seminar on Stability Problems for Stochastic Models and V International Workshop "Applied Problems in Theory of Probabilities and Mathematical Statistics related to modeling of information systems". Book of Abstracts. 2011. Pp. 82–83.

Examples

```
StabC(s=2,p=c(.5,.5), method="less-memory")
# returns the constant for a 2-server model,
# where a task needs 1 or 2 cores with equal probability
```

ToSWF *Convertor from a dataset to Standart Workload Format*

Description

This is a convertor from a correctly built dataset to a Standart Workload Format used to share the logfiles of High Performance Clusters

Usage

```
ToSWF(T, S, N, D, filename="output.swf")
```

Arguments

T	Interarrival times of tasks (a vector)
S	Service times of tasks (a vector)
N	Number of cores each task needs (a vector)
D	The delays of tasks in a queue (a vector)
filename	The file to store the converted workload (output.swf by default)

Details

The Standart Workload Format is a single format to store and exchange high performance cluster logs, that is used in Parallel Workload Archive. See references for current standard. The SWF format may contain additional data, but in this package only the 1st to 5th fields are used. One may also need to manually fill in the header of the file in order to completely prepare the resulting SWF file.

Value

Nothing is returned, but a file is created in the current working directory (with default name output.swf) containing the converted data.

Author(s)

Alexander Rumyantsev (Institute of Applied Mathematical Research, Karelian Research Centre, RAS)

References

Feitelson, D.G. and Tsafrir, D. and Krakov D. 2012 Experience with the Parallel Workloads Archive. Technical Report 2012-6, School of Computer Science and Engineering, the Hebrew University April, 2012, Jerusalem, Israel

http://www.cs.huji.ac.il/labs/parallel/workload/swf.html

See Also

FromSWF,DataToSWF

Examples

```
##---- Should be DIRECTLY executable !! ----
##-- ==>  Define data, use random,
##--or do  help(data=index)  for the standard data sets.
data(HPC_KRC)
ToSWF(HPC_KRC$interarrival, HPC_KRC$service, HPC_KRC$cores_requested, HPC_KRC$delay)
```

Wld *Workload of a High Performance Cluster model*

Description

This function computes the Kiefer-Wolfowitz modified vector for a HPC model. This vector contains the work left on each of 'm' servers of a cluster for the time of the arival of a task. Two methods are available, one for the case of concurrent server release (all the servers end a single task simultaneously), other for independent release (service times on each server are independent).

Usage

```
Wld(T, S, N, m, method = "concurrent")
```

Arguments

T	Interarrival times of tasks
S	Service times of tasks (a vector of length n, or a matrix nrows=n, ncols='m').
N	Number of cores each task needs
m	Number of cores/servers for a HPC
method	Independent or concurrent

Value

A dataset is returned, containing 'delay' as a vector of delays exhibited by each task, 'total_cores' as the total busy CPUs in time of arrival of each task, and 'workload' as total work left at each CPU.

Author(s)

Alexander Rumyantsev (Institute of Applied Mathematical Research, Karelian Research Centre, RAS)

References

E.V. Morozov, A.Rumyantsev. Stability analysis of a multiprocessor model describing a high performance cluster. XXIX International Seminar on Stability Problems for Stochastic Models and V International Workshop "Applied Problems in Theory of Probabilities and Mathematical Statistics related to modeling of information systems". Book of Abstracts. 2011. Pp. 82–83.

Examples

```
Wld(T=rexp(1000,1), S=rexp(1000,1), round(runif(1000,1,10)), 10)
# returns the workload, delay and total cpus used
# for a cluster with 10 CPUs and random exponential times
```

X *Dataset with raw workload data from HPDC KRC RAS*

Description

Source data for workload of HPC of HPDC KRC RAS. More usable dataset is HPC_KRC. This are raw times in sec. since 1 January 1970, for tasks arrival times, start of execution times and end times.

Usage

```
data(X)
```

Format

The format is: num [1:8499, 1:3] 1.24e+09 1.24e+09 1.24e+09 1.24e+09 1.24e+09 ...

Source

http://cluster.krc.karelia.ru

References

http://cluster.krc.karelia.ru

Examples

```
data(X)
```

Index

∗Topic **\textasciitildecluster**
  ApproxC, 3
  DMC, 5
  StabC, 9
  Wld, 11
∗Topic **\textasciitildecorrelation**
  DMC, 5
∗Topic **\textasciitildeformat**
  DataToSWF, 4
  FromSWF, 6
  ToSWF, 10
∗Topic **\textasciitildestability**
  ApproxC, 3
  StabC, 9
∗Topic **\textasciitildestandard**
  DataToSWF, 4
  FromSWF, 6
  ToSWF, 10
∗Topic **\textasciitildeworkload**
  DataToSWF, 4
  FromSWF, 6
  ToSWF, 10
  Wld, 11
∗Topic **datasets**
  HPC_KRC, 7
  HPC_KRC2, 8
  X, 12
∗Topic **package**
  hpcwld-package, 2

ApproxC, 3

DataToSWF, 4, *6*, *11*
DMC, 5

FromSWF, *5*, 6, *11*

HPC_KRC, 7
HPC_KRC2, 8
hpcwld *(*hpcwld-package*)*, 2

hpcwld-package, 2

StabC, 9

ToSWF, *5*, *6*, 10

Wld, 11

X, 12