

Package ‘HDF5Array’

April 12, 2018

Title HDF5 back end for DelayedArray objects

Description An array-like container for convenient access and manipulation of HDF5 datasets. Supports delayed operations and block processing.

Version 1.6.0

Encoding UTF-8

Author Hervé Pagès

Maintainer Hervé Pagès <hpages@fredhutch.org>

biocViews Infrastructure, DataRepresentation, Sequencing, Annotation, Coverage, GenomeAnnotation

Depends R (>= 3.4), methods, DelayedArray (>= 0.3.18), rhdf5

Imports utils, tools, BiocGenerics, S4Vectors, IRanges

Suggests h5vcData, SummarizedExperiment (>= 1.5.6), GenomicRanges, BiocStyle

License Artistic-2.0

Collate utils.R HDF5Array-class.R dump-management.R writeHDF5Array.R
zzz.R

NeedsCompilation no

R topics documented:

HDF5-dump-management	1
HDF5Array-class	5
writeHDF5Array	7
Index	9

HDF5-dump-management *HDF5 dump management*

Description

A set of utilities to control the location of automatically created HDF5 datasets.

Usage

```

setHDF5DumpDir(dir)
setHDF5DumpFile(file)
setHDF5DumpName(name)
setHDF5DumpCompressionLevel(level=6L)

getHDF5DumpDir()
getHDF5DumpFile(for.use=FALSE)
getHDF5DumpName(for.use=FALSE)
getHDF5DumpCompressionLevel()

lsHDF5DumpFile()

showHDF5DumpLog()

## For developers:
getHDF5DumpChunkDim(dim, type, ratio=75)
appendDatasetCreationToHDF5DumpLog(file, name, dim, type, chunk_dim, level)

```

Arguments

dir	The path (as a single string) to the current <i>HDF5 dump directory</i> , that is, to the (new or existing) directory where <i>HDF5 dump files</i> with automatic names will be created. This is ignored if the user specified an <i>HDF5 dump file</i> with <code>setHDF5DumpFile</code> . If <code>dir</code> is missing, then the <i>HDF5 dump directory</i> is set back to its default value i.e. to some directory under <code>tempdir()</code> (call <code>getHDF5DumpDir()</code> to get the exact path).
file	For <code>setHDF5DumpFile</code> : The path (as a single string) to the current <i>HDF5 dump file</i> , that is, to the (new or existing) HDF5 file where the <i>next automatic HDF5 datasets</i> will be written. If <code>file</code> is missing, then a new file with an automatic name will be created (in <code>getHDF5DumpDir()</code>) and used for each new dataset. For <code>appendDatasetCreationToHDF5DumpLog</code> : See the Note TO DEVELOPERS below.
name	For <code>setHDF5DumpName</code> : The name of the <i>next automatic HDF5 dataset</i> to be written to the current <i>HDF5 dump file</i> . For <code>appendDatasetCreationToHDF5DumpLog</code> : See the Note TO DEVELOPERS below.
level	For <code>setHDF5DumpCompressionLevel</code> : The compression level to use for writing <i>automatic HDF5 datasets</i> to disk. See the <code>level</code> argument in <code>?rhdf5::h5createDataset</code> (in the rhdf5 package) for more information about this. For <code>appendDatasetCreationToHDF5DumpLog</code> : See the Note TO DEVELOPERS below.
for.use	Whether the returned file or dataset name is for use by the caller or not. See below for the details.
dim	The dimensions of the HDF5 dataset to be written to disk, that is, an integer vector of length one or more giving the maximal indices in each dimension. See the <code>dims</code> argument in <code>?rhdf5::h5createDataset</code> (in the rhdf5 package) for more information about this.
type	The type (a.k.a. storage mode) of the data to be written to disk. Can be obtained with <code>type()</code> on an array-like object (which is equivalent to <code>storage.mode()</code> or

typeof() on an ordinary array). This is typically what an application writing datasets to the *HDF5 dump* should pass to the `storage.mode` argument of its call to `rhdf5::h5createDataset`. See the Note TO DEVELOPERS below for more information.

`ratio` The number of chunks per block. By default, a ratio of 75 is used.

`chunk_dim` The dimensions of the chunks.

Details

Calling `getHDF5DumpFile()` and `getHDF5DumpName()` with no argument should be *informative* only i.e. it's a mean for the user to know where the *next automatic HDF5 dataset* will be written. Since a given file/name combination can be used only once, the user should be careful to not use that combination to explicitly create an HDF5 dataset because that would get in the way of the creation of the *next automatic HDF5 dataset*. See the Note TO DEVELOPERS below if you actually need to use this file/name combination.

`lsHDF5DumpFile()` is a just convenience wrapper for `rhdf5::h5ls(getHDF5DumpFile())`.

Value

`getHDF5DumpDir` returns the absolute path to the directory where *HDF5 dump files* with automatic names will be created. Only meaningful if the user did NOT specify an *HDF5 dump file* with `setHDF5DumpFile`.

`getHDF5DumpFile` returns the absolute path to the HDF5 file where the *next automatic HDF5 dataset* will be written.

`getHDF5DumpName` returns the name of the *next automatic HDF5 dataset*.

`getHDF5DumpCompressionLevel` returns the compression level currently used for writing *automatic HDF5 datasets* to disk.

`showHDF5DumpLog` returns the dump log in an invisible data frame.

`getHDF5DumpChunkDim` returns the dimension of the chunks for the specified ratio. By default a ratio of 75 is used i.e. the *automatic HDF5 dataset* are written to disk using 75 chunks per block.

Note

TO DEVELOPERS:

If your application needs to write its own dataset to the *HDF5 dump* then it should:

1. Get a file/name combination by calling `getHDF5DumpFile(for.use=TRUE)` and `getHDF5DumpName(for.use=TRUE)`

OPTIONAL Call `getHDF5DumpChunkDim(dim, type)` (possibly with a non-default ratio) to get reasonable chunk dimensions to use for writing the data to disk. Or choose your own chunk dimensions.

2. Add an entry to the dump log by calling `appendDatasetCreationToHDF5DumpLog`. Typically, this should be done right after creating the dataset (e.g. with `rhdf5::h5createDataset`) and before starting to write the data to disk. The values passed to `appendDatasetCreationToHDF5DumpLog` via the `file`, `name`, `dim`, `type`, `chunk_dim`, and `level` arguments should be those that were passed to `rhdf5::h5createDataset` via the `file`, `dataset`, `dims`, `storage.mode`, `chunk`, and `level` arguments, respectively. Note that `appendDatasetCreationToHDF5DumpLog` uses a lock mechanism so is safe to use in the context of parallel execution.

This is actually what the coercion method to `HDF5Array` does internally.

See Also

- [writeHDF5Array](#) for writing an array-like object to an HDF5 file.
- [HDF5Array](#) objects.
- The [h5ls](#) function in the **rhdf5** package, on which `lsHDF5DumpFile` is based.
- [type](#) in the **DelayedArray** package.

Examples

```

getHDF5DumpDir()
getHDF5DumpFile()

## Use setHDF5DumpFile() to change the current HDF5 dump file.
## If the specified file exists, then it must be in HDF5 format or
## an error will be raised. If it doesn't exist, then it will be
## created.
setHDF5DumpFile("path/to/some/HDF5/file")

lsHDF5DumpFile()

a <- array(1:600, c(150, 4))
A <- as(a, "HDF5Array")
lsHDF5DumpFile()
A

b <- array(runif(6000), c(4, 2, 150))
B <- as(b, "HDF5Array")
lsHDF5DumpFile()
B

C <- (log(2 * A + 0.88) - 5)^3 * t(drop(B[, 1, ]))
as(C, "HDF5Array") # realize C on disk
lsHDF5DumpFile()

## Matrix multiplication is not delayed: the output matrix is realized
## block by block. The current "realization backend" controls where
## realization happens e.g. in memory if set to NULL or in an HDF5 file
## if set to "HDF5Array". See '?realize' in the DelayedArray package for
## more information about "realization backends".
setRealizationBackend("HDF5Array")
m <- matrix(runif(20), nrow=4)
P <- C %*% m
lsHDF5DumpFile()

## See all the HDF5 datasets created in the current session so far:
showHDF5DumpLog()

## Wrap the call in suppressMessages() if you are only interested in the
## data frame version of the dump log:
dump_log <- suppressMessages(showHDF5DumpLog())
dump_log

```

Description

We provide 2 classes for representing an (on-disk) HDF5 dataset as an array-like object in R:

- **HDF5Array**: A high-level class **HDF5Array** that extends [DelayedArray](#). All the operations available on [DelayedArray](#) objects work on **HDF5Array** objects.
- **HDF5ArraySeed**: A low-level class for pointing to an HDF5 dataset. No operation can be performed directly on an **HDF5ArraySeed** object. It needs to be wrapped in a [DelayedArray](#) or **HDF5Array** object first. An **HDF5Array** object is just an **HDF5ArraySeed** object wrapped in a [DelayedArray](#) object.

Usage

```
## Constructor functions:  
HDF5Array(file, name, type=NA)  
HDF5ArraySeed(file, name, type=NA)
```

Arguments

file	The path (as a single character string) to the HDF5 file where the dataset is located.
name	The name of the dataset in the HDF5 file.
type	NA or the <i>R atomic type</i> (specified as a single string) corresponding to the type of the HDF5 dataset.

Value

An **HDF5Array** object for **HDF5Array()**.

An **HDF5ArraySeed** object for **HDF5ArraySeed()**.

See Also

- [DelayedArray](#) objects.
- [DelayedArray-utils](#) for common operations on [DelayedArray](#) objects.
- [writeHDF5Array](#) for writing an array-like object to an HDF5 file.
- [HDF5-dump-management](#) for controlling the location of automatically created HDF5 datasets.
- [saveHDF5SummarizedExperiment](#) and [loadHDF5SummarizedExperiment](#) in the **SummarizedExperiment** package for saving/loading a HDF5-based **SummarizedExperiment** object to/from disk.
- [h5ls](#) in the **rhdf5** package.
- The **rhdf5** package on top of which **HDF5Array** objects are implemented.
- [array](#) objects in base R.

Examples

```

## -----
## CONSTRUCTION
## -----
library(rhdf5)
library(h5vcData)

tally_file <- system.file("extdata", "example.tally.hfs5",
                          package="h5vcData")
h5ls(tally_file)

## Pick up "Coverages" dataset for Human chromosome 16:
cov0 <- HDF5Array(tally_file, "/ExampleStudy/16/Coverages")
cov0

## -----
## dim/dimnames
## -----
dim(cov0)

dimnames(cov0)
dimnames(cov0) <- list(paste0("s", 1:6), c("+", "-"))
dimnames(cov0)

## -----
## SLICING (A.K.A. SUBSETTING)
## -----
cov1 <- drop(cov0[ , , 29000001:29000007])
cov1

dim(cov1)
as.array(cov1)
stopifnot(identical(dim(as.array(cov1)), dim(cov1)))
stopifnot(identical(dimnames(as.array(cov1)), dimnames(cov1)))

cov2 <- drop(cov0[ , "+", 29000001:29000007])
cov2
as.matrix(cov2)

## -----
## SummarizedExperiment OBJECTS WITH DELAYED ASSAYS
## -----

## DelayedArray objects can be used inside a SummarizedExperiment object
## to hold the assay data and to delay operations on them.

library(SummarizedExperiment)

pcov <- drop(cov0[ , 1, ]) # coverage on plus strand
mcov <- drop(cov0[ , 2, ]) # coverage on minus strand

nrow(pcov) # nb of samples
ncol(pcov) # length of Human chromosome 16

## The convention for a SummarizedExperiment object is to have 1 column
## per sample so first we need to transpose 'pcov' and 'mcov':

```

```

pcov <- t(pcov)
mcof <- t(mcov)
se <- SummarizedExperiment(list(pcov=pcov, mcof=mcof))
se
stopifnot(validObject(se, complete=TRUE))

## A GPos object can be used to represent the genomic positions along
## the dataset:
gpos <- GPos(GRanges("16", IRanges(1, nrow(se))))
gpos
rowRanges(se) <- gpos
se
stopifnot(validObject(se))
assays(se)$pcov
assays(se)$mcof

```

writeHDF5Array	<i>Write an array-like object to an HDF5 file</i>
----------------	---

Description

A function for writing an array-like object to an HDF5 file.

Usage

```
writeHDF5Array(x, file=NULL, name=NULL, chunk_dim=NULL, level=NULL,
              verbose=FALSE)
```

Arguments

x	The array-like object to write to an HDF5 file. If x is a DelayedArray object, <code>writeHDF5Array</code> <i>realizes</i> it on disk, that is, all the delayed operations carried by the object are executed while the object is written to disk. See "On-disk realization of a DelayedArray object as an HDF5 dataset" section below for more information.
file	NULL or the path (as a single string) to the (new or existing) HDF5 file where to write the dataset. If NULL, then the dataset will be written to the current <i>HDF5 dump file</i> i.e. the path returned by getHDF5DumpFile will be used.
name	NULL or the name of the HDF5 dataset to write. If NULL, then the name returned by getHDF5DumpName will be used.
chunk_dim	The dimensions of the chunks to use for writing the data to disk. By default, <code>getHDF5DumpChunkDim(dim(x), type(x))</code> is used. See ?getHDF5DumpChunkDim for more information.
level	The compression level to use for writing the data to disk. By default, <code>getHDF5DumpCompressionLevel</code> is used. See ?getHDF5DumpCompressionLevel for more information.
verbose	Set to TRUE to make the function display progress.

Details

Please note that, depending on the size of the data to write to disk and the performance of the disk, `writeHDF5Array` can take a long time to complete. Use `verbose=TRUE` to see its progress.

Use [setHDF5DumpFile](#) and [setHDF5DumpName](#) to control the location of automatically created HDF5 datasets.

Value

An HDF5Array object pointing to the newly written HDF5 dataset on disk.

On-disk realization of a DelayedArray object as an HDF5 dataset

When passed a [DelayedArray](#) object, `writeHDF5Array` *realizes* it on disk, that is, all the delayed operations carried by the object are executed on-the-fly while the object is written to disk. This uses a block-processing strategy so that the full object is not realized at once in memory. Instead the object is processed block by block i.e. the blocks are realized in memory and written to disk one at a time.

In other words, `writeHDF5Array(x, ...)` is semantically equivalent to `writeHDF5Array(as.array(x), ...)`, except that `as.array(x)` is not called because this would realize the full object at once in memory.

See [?DelayedArray](#) for general information about [DelayedArray](#) objects.

See Also

- [saveHDF5SummarizedExperiment](#) and [loadHDF5SummarizedExperiment](#) in the **SummarizedExperiment** package for saving/loading a HDF5-based SummarizedExperiment object to/from disk.
- [HDF5-dump-management](#) for controlling the location of automatically created HDF5 datasets.
- [HDF5Array](#) objects.
- [DelayedArray](#) objects.
- [DelayedArray-utils](#) for common operations on DelayedArray objects.

Examples

```
library(rhdf5)
library(h5vcData)

tally_file <- system.file("extdata", "example.tally.hfs5",
                          package="h5vcData")
h5ls(tally_file)

cov0 <- HDF5Array(tally_file, "/ExampleStudy/16/Coverages")

cov1 <- drop(cov0[ , , 29000001:29000007])

out_file <- tempfile()
writeHDF5Array(cov1, out_file, "cov1")
h5ls(out_file)
```


Index

*Topic **classes**

HDF5Array-class, 5

*Topic **methods**

HDF5-dump-management, 1

HDF5Array-class, 5

writeHDF5Array, 7

appendDatasetCreationToHDF5DumpLog
(HDF5-dump-management), 1

array, 5

chunk_dim, HDF5RealizationSink-method
(writeHDF5Array), 7

class:HDF5Array (HDF5Array-class), 5

class:HDF5ArraySeed (HDF5Array-class), 5

class:HDF5Matrix (HDF5Array-class), 5

class:HDF5RealizationSink
(writeHDF5Array), 7

coerce, ANY, HDF5Array-method
(writeHDF5Array), 7

coerce, ANY, HDF5Matrix-method
(HDF5Array-class), 5

coerce, DelayedArray, HDF5Array-method
(writeHDF5Array), 7

coerce, DelayedMatrix, HDF5Matrix-method
(writeHDF5Array), 7

coerce, HDF5Array, HDF5Matrix-method
(HDF5Array-class), 5

coerce, HDF5RealizationSink, DelayedArray-method
(writeHDF5Array), 7

coerce, HDF5RealizationSink, HDF5Array-method
(writeHDF5Array), 7

coerce, HDF5RealizationSink, HDF5ArraySeed-method
(writeHDF5Array), 7

DelayedArray, 5, 7, 8

DelayedArray, HDF5ArraySeed-method
(HDF5Array-class), 5

DelayedArray-utils, 5, 8

dim, HDF5ArraySeed-method
(HDF5Array-class), 5

dimnames, HDF5RealizationSink-method
(writeHDF5Array), 7

dump-management (HDF5-dump-management),
1

getHDF5DumpChunkDim, 7

getHDF5DumpChunkDim
(HDF5-dump-management), 1

getHDF5DumpCompressionLevel, 7

getHDF5DumpCompressionLevel
(HDF5-dump-management), 1

getHDF5DumpDir (HDF5-dump-management), 1

getHDF5DumpFile, 7

getHDF5DumpFile (HDF5-dump-management),
1

getHDF5DumpName, 7

getHDF5DumpName (HDF5-dump-management),
1

h5createDataset, 2, 3

h5ls, 3–5

HDF5-dump-management, 1, 5, 8

HDF5Array, 3, 4, 8

HDF5Array (HDF5Array-class), 5

HDF5Array-class, 5

HDF5ArraySeed (HDF5Array-class), 5

HDF5ArraySeed-class (HDF5Array-class), 5

HDF5Matrix (HDF5Array-class), 5

HDF5Matrix-class (HDF5Array-class), 5

HDF5RealizationSink (writeHDF5Array), 7

HDF5RealizationSink-class
(writeHDF5Array), 7

loadHDF5SummarizedExperiment, 5, 8

lsHDF5DumpFile (HDF5-dump-management), 1

matrixClass, HDF5Array-method
(HDF5Array-class), 5

rhdf5, 5

saveHDF5SummarizedExperiment, 5, 8

setHDF5DumpCompressionLevel
(HDF5-dump-management), 1

setHDF5DumpDir (HDF5-dump-management), 1

setHDF5DumpFile, 7

setHDF5DumpFile (HDF5-dump-management),
1

setHDF5DumpName, 7

setHDF5DumpName (HDF5-dump-management),
1

showHDF5DumpLog (HDF5-dump-management),
1

subset_seed_as_array, HDF5ArraySeed-method
(HDF5Array-class), 5

type, 2, 4

write_block_to_sink, HDF5RealizationSink-method
(writeHDF5Array), 7

writeHDF5Array, 4, 5, 7

writeHDF5Dataset (writeHDF5Array), 7