

Package: beachmat (via r-universe)

October 24, 2026

Version 2.29.3

Date 2026-09-23

Title Compiling Bioconductor to Handle Each Matrix Type

Imports methods, DelayedArray (>= 0.27.2), SparseArray, BiocGenerics, Matrix, Rcpp

Suggests testthat, BiocStyle, knitr, rmarkdown, rcmdcheck, BiocParallel, HDF5Array, beachmat.hdf5

LinkingTo Rcpp, assorthead (>= 1.5.4)

biocViews DataRepresentation, DataImport, Infrastructure

Description Provides a consistent C++ class interface for reading from a variety of commonly used matrix types. Ordinary matrices and several sparse/dense Matrix classes are directly supported, along with a subset of the delayed operations implemented in the DelayedArray package. All other matrix-like objects are supported by calling back into R.

License GPL-3

NeedsCompilation yes

VignetteBuilder knitr

SystemRequirements C++17

URL <https://github.com/tatami-inc/beachmat>

BugReports <https://github.com/tatami-inc/beachmat/issues>

Encoding UTF-8

Config/roxygen2/version 8.1.0

Config/pak/sysreqs zlib1g-dev

Repository <https://bioc.r-universe.dev>

Date/Publication 2026-09-23 05:33:08 UTC

RemoteUrl <https://github.com/bioc/beachmat>

RemoteRef HEAD

RemoteSha 7fe7dc234ea2fe69bf8f4c711525428539033652

Contents

checkMemoryCache	2
colBlockApply	3
getExecutor	6
initializeCpp	6
realizeFileBackedMatrix	7
tatami-utils	8
toCsparse	12
whichNonZero	13
Index	14

checkMemoryCache	<i>Check the in-memory cache for matrix instances</i>
------------------	---

Description

Check the in-memory cache for a pre-existing initialized C++ object, and initialize it if it does not exist. This is typically used in [initializeCpp](#) methods of file-backed representations to avoid redundant reads of the entire matrix.

Usage

```
flushMemoryCache()
```

```
checkMemoryCache(namespace, key, fun)
```

Arguments

namespace	String containing the namespace, typically the name of the package implementing the method.
key	String containing the key for a specific matrix instance.
fun	Function that accepts no arguments and returns an external pointer like those returned by initializeCpp .

Details

For representations where data extraction is costly (e.g., from file), [initializeCpp](#) methods may consider realizing the entire matrix into memory. This effectively pays a one-time up-front cost to improve efficiency for downstream operations that pass through the matrix multiple times.

If such an option is enabled, [initializeCpp](#) methods are expected to cache the in-memory instance using [checkMemoryCache](#). This ensures that all subsequent calls to the same [initializeCpp](#) method will return the same instance, avoiding redundant memory loads when the same matrix is used in multiple functions.

Of course, this process comes at the expense of increased memory usage. If too many instances are stored in the cache, they can be cleared from memory using the [flushMemoryCache](#) function.

Value

For `checkMemoryCache`, the output of `fun` (possibly from an existing cache) is returned.

For `flushMemoryCache`, all existing cached objects are removed and `NULL` is invisibly returned.

Author(s)

Aaron Lun

See Also

The `hdf5.realize=` and `tiledb.realize=` options in the `initializeCpp` methods of **beachmat.hdf5** and **beachmat.tiledb**.

Examples

```
# Mocking up a class with some kind of uniquely identifying aspect.
setClass("UnknownMatrix", slots=c(contents="dgCMatrix", uuid="character"))
X <- new("UnknownMatrix",
  contents=Matrix::rsparsematrix(10, 10, 0.1),
  uuid=as.character(sample(1e8, 1)))

# Defining our initialization method.
setMethod("initializeCpp", "UnknownMatrix", function(x, ..., memorize=FALSE) {
  if (memorize) {
    checkMemoryCache("my_package", x@uuid, function() initializeCpp(x@contents))
  } else {
    initializeCpp(x@contents)
  }
})

# Same pointer is returned multiple times.
initializeCpp(X, memorize=TRUE)
initializeCpp(X, memorize=TRUE)

# Flushing the cache.
flushMemoryCache()
```

colBlockApply

Apply over blocks of columns or rows

Description

Apply a function over blocks of columns or rows using **DelayedArray**'s block processing mechanism.

Usage

```
colBlockApply(
  x,
  FUN,
  ...,
  grid = NULL,
  coerce.sparse = TRUE,
  BPPARAM = getAutoBPPARAM()
)

rowBlockApply(
  x,
  FUN,
  ...,
  grid = NULL,
  coerce.sparse = TRUE,
  BPPARAM = getAutoBPPARAM()
)
```

Arguments

x	A matrix-like object to be split into blocks and looped over. This can be of any class that respects the matrix contract.
FUN	A function that operates on columns or rows in x, for colBlockApply and rowBlockApply respectively. Ordinary matrices, CsparseMatrix or SparseMatrix objects may be passed as the first argument.
...	Further arguments to pass to FUN.
grid	An ArrayGrid object specifying how x should be split into blocks. For colBlockApply and rowBlockApply, blocks should consist of consecutive columns and rows, respectively. Alternatively, this can be set to TRUE or FALSE, see Details.
coerce.sparse	Logical scalar indicating whether blocks of a sparse DelayedMatrix x should be automatically coerced into CsparseMatrix objects.
BPPARAM	A BiocParallelParam object from the BiocParallel package, specifying how parallelization should be performed across blocks.

Details

This is a wrapper around [blockApply](#) that is dedicated to looping across rows or columns of x. The aim is to provide a simpler interface for the common task of [applying](#) across a matrix, along with a few modifications to improve efficiency for parallel processing and for natively supported x.

Note that the fragmentation of x into blocks is not easily predictable, meaning that FUN should be capable of operating on each row/column independently. Users can retrieve the current location of each block of x by calling [currentViewport](#) inside FUN.

If grid is not explicitly set to an [ArrayGrid](#) object, it can take several values:

- If TRUE, the function will choose a grid that (i) respects the memory limits in [getAutoBlockSize](#) and (ii) fragments x into sufficiently fine chunks that every worker in BPPARAM gets to do something. If FUN might make large allocations, this mode should be used to constrain memory usage.
- The default grid=NULL is very similar to TRUE except that that memory limits are ignored when x is of any type that can be passed directly to FUN. This avoids unnecessary copies of x and is best used when FUN itself does not make large allocations.
- If FALSE, the function will choose a grid that covers the entire x. This is provided for completeness and is only really useful for debugging.

The default of `coerce.sparse=TRUE` will generate `dgCMatrix` objects during block processing of a sparse `DelayedMatrix` x. This is convenient as it avoids the need for FUN to specially handle [SparseMatrix](#) objects from the **SparseArray** package. If the coercion is not desired (e.g., to preserve integer values in x), it can be disabled with `coerce.sparse=FALSE`.

Value

A list of length equal to the number of blocks, where each entry is the output of FUN for the results of processing each the rows/columns in the corresponding block.

See Also

[blockApply](#), for the original **DelayedArray** implementation.

[toCsparse](#), to convert `SparseMatrix` objects to `CsparseMatrix` objects prior to further processing in FUN.

Examples

```
x <- matrix(runif(10000), ncol=10)
str(colBlockApply(x, colSums))
str(rowBlockApply(x, rowSums))

library(Matrix)
y <- rsparsematrix(10000, 10000, density=0.01)
str(colBlockApply(y, colSums))
str(rowBlockApply(y, rowSums))

library(DelayedArray)
z <- DelayedArray(y) + 1
str(colBlockApply(z, colSums))
str(rowBlockApply(z, rowSums))

# We can also force multiple blocks:
library(BiocParallel)
BPPARAM <- SnowParam(2)
str(colBlockApply(x, colSums, BPPARAM=BPPARAM))
str(rowBlockApply(x, rowSums, BPPARAM=BPPARAM))
```

getExecutor	<i>Get the parallel executor</i>
-------------	----------------------------------

Description

Get the executor object for safe execution of R code in parallel sections. This should be set by `Rtatami::set_executor()` in the `.onLoad` function of downstream packages.

Usage

```
getExecutor()
```

Value

An external pointer to be passed to `Rtatami::set_executor`.

Author(s)

Aaron Lun

Examples

```
getExecutor()
```

initializeCpp	<i>Initialize matrix in C++ memory space</i>
---------------	--

Description

Initialize a **tatami** matrix object in C++ memory space from an abstract numeric R matrix. This object simply references the R memory space and avoids making any copies of its own, so it can be cheaply re-created when needed inside each function.

Usage

```
initializeCpp(x, ...)
```

Arguments

<code>x</code>	A matrix-like object, typically from the Matrix or DelayedArray packages. Alternatively, an external pointer from a previous call to <code>initializeCpp</code> , which is returned without modification.
<code>...</code>	Further arguments used by specific methods, such as:

- `.check.na`, a boolean indicating whether to check for NA values in integer and logical matrices. If TRUE (the default), any NAs are cast to their double-precision equivalents when reading from the tatami matrix. This can be set to FALSE to improve performance if the caller knows that `x` does not contain NAs.
- `.unknown.action`, a string specifying the action to take upon encountering a matrix with no known **tatami** representation. This should be one of "message", "warn", "error", or "none".

If a `initializeCpp` method accepts additional arguments, the names of those argument should generally be prefixed by the matrix type to avoid conflicts between different methods. For example, `hdf5.realize` can be used in **beach-mat.hdf5** to load a HDF5-backed matrix into memory.

Details

Do not attempt to serialize the return value; it contains a pointer to external memory, and will not be valid after a save/load cycle. Users should not be exposed to the returned pointers; rather, developers should call `initializeCpp` at the start to obtain a C++ object for further processing. The initialization process should be cheap so there is no downside from just recreating the object within each function body.

Value

An external pointer to a C++ object containing a tatami matrix.

Examples

```
# Mocking up a count matrix:
x <- Matrix::rsparsematrix(1000, 100, 0.1)
y <- round(abs(x))

stuff <- initializeCpp(y)
stuff
```

```
realizeFileBackedMatrix
```

Realize a file-backed DelayedMatrix

Description

Realize a file-backed `DelayedMatrix` into its corresponding in-memory format.

Usage

```
realizeFileBackedMatrix(x)

isFileBackedMatrix(x)
```

Arguments

`x` A [DelayedMatrix](#) object.

Details

A file-backed matrix representation is recognized based on whether it has a [path](#) method for any one of its seeds. If so, and the "beachmat.realizeFileBackedMatrix" option is not FALSE, we will load it into memory. This is intended for DelayedMatrix objects that have already been subsetting (e.g., to highly variable genes), which can be feasibly loaded into memory for rapid calculations.

Value

For `realizeFileBackedMatrix`, an ordinary matrix or a `dgCMatrix`, depending on whether `is_sparse(x)`.
For `isFileBackedMatrix`, a logical scalar indicating whether `x` has file-backed components.

Author(s)

Aaron Lun

Examples

```
mat <- matrix(rnorm(50), ncol=5)
realizeFileBackedMatrix(mat) # no effect

library(HDF5Array)
mat2 <- as(mat, "HDF5Array")
realizeFileBackedMatrix(mat2) # realized into memory
```

tatami-utils

Tatami utilities

Description

Utility functions that directly operate on the pointers produced by [initializeCpp](#). Some of these are used internally by `initializeCpp` methods operating on **DelayedArray** classes.

Usage

```
tatami.bind(xs, by.row)

tatami.transpose(x)

tatami.subset(x, subset, by.row)

tatami.arith(x, op, val, by.row, right)

tatami.compare(x, op, val, by.row, right)
```

```

tatami.logic(x, op, val, by.row)

tatami.round(x)

tatami.log(x, base)

tatami.math(x, op)

tatami.not(x)

tatami.binary(x, y, op)

tatami.dim(x)

tatami.get(x, i, row)

tatami.is.sparse(x)

tatami.prefer.rows(x)

tatami.realize(x, num.threads)

tatami.multiply(x, val, right, num.threads)

tatami.sums(x, row, num.threads)

tatami.variances(x, row, num.threads)

tatami.sums.by.group(x, group, num.groups, row, num.threads)

tatami.medians(x, row, num.threads)

tatami.nan.counts(x, row, num.threads)

```

Arguments

<code>xs</code>	A list of pointers produced by initializeCpp . All matrices should have the same number of rows (if <code>by.row=FALSE</code>) or columns (otherwise).
<code>by.row</code>	Logical scalar indicating whether to apply the operation on the rows. <ul style="list-style-type: none"> For <code>tatami.bind</code>, this will combine the matrices by rows, i.e., the output matrix has a number of rows equal to the sum of the number of rows in <code>xs</code>. For <code>tatami.subset</code>, this will subset the matrix by row. For <code>tatami.arith</code>, <code>tatami.compare</code> and <code>tatami.logic</code> with a vector <code>val</code>, the vector should have length equal to the number of rows <code>k</code>.
<code>x</code>	A pointer produced by initializeCpp .
<code>subset</code>	Integer vector containing the subset of interest. These should be 1-based row or column indices depending on <code>by.row</code> .

op	<p>String specifying the operation to perform.</p> <ul style="list-style-type: none"> • For <code>tatami.arith</code>, this should be one of the operations in Arith. • For <code>tatami.compare</code>, this should be one of the operations in Compare. • For <code>tatami.logic</code>, this should be one of the operations in Logic. • For <code>tatami.math</code>, this should be one of the operations in Math. • For <code>tatami.binary</code>, this may be any operation in Arith, Compare or Logic.
val	<p>For <code>tatami.arith</code>, <code>tatami.compare</code> and <code>tatami.logic</code>, the value to be used in the operation specified by <code>op</code>. This may be a:</p> <ul style="list-style-type: none"> • Numeric scalar, which is used in the operation for all entries of the matrix. • Numeric vector of length equal to the number of rows, where each value is used in the operation with the corresponding row when <code>by.row=TRUE</code>. • Numeric vector of length equal to the number of column, where each value is used with the corresponding column when <code>by.row=FALSE</code>. <p>For <code>tatami.multiply</code>, the value to be used in the matrix multiplication. This may be a:</p> <ul style="list-style-type: none"> • Numeric vector of length equal to the number of columns of <code>x</code> (if <code>right=FALSE</code>) or rows (otherwise). • Numeric matrix with number of rows equal to the number of columns of <code>x</code> (if <code>right=FALSE</code>) or rows (otherwise). • Pointer produced by initializeCpp, referencing a matrix with number of rows equal to the number of columns of <code>x</code> (if <code>right=FALSE</code>) or rows (otherwise).
right	<p>For <code>tatami.arith</code> and <code>tatami.compare</code>, a logical scalar indicating that <code>val</code> is on the right-hand side of the operation.</p> <p>For <code>tatami.multiply</code>, a logical scalar indicating that <code>val</code> is on the right-hand side of the multiplication.</p>
base	Numeric scalar specifying the base of the log-transformation.
y	A pointer produced by initializeCpp , referencing a matrix of the same dimensions as <code>x</code> .
i	Integer scalar containing the 1-based index of the row (for <code>row=TRUE</code>) or column (otherwise) of interest. This should be in <code>[1, D]</code> where <code>D</code> is the total number of rows or columns, respectively, in <code>x</code> .
row	<p>For <code>tatami.get</code>, a boolean indicating whether to extract the <code>i</code>-th row. If <code>FALSE</code>, the <code>i</code>-th column is extracted instead.</p> <p>For <code>tatami.sums</code>, <code>tatami.sums.by.group</code>, <code>tatami.medians</code>, etc., a boolean indicating whether to compute row-wise statistics. If <code>FALSE</code>, column-wise statistics are computed instead.</p>
num.threads	Integer scalar specifying the number of threads to use.
group	Integer vector of length equal to the number of columns (if <code>row = TRUE</code>) or rows (otherwise), containing the group assignment for each column and row, respectively. Assignments should lie in <code>[1, num.groups]</code> .
num.groups	Integer specifying the total number of unique groups in <code>group</code> .

Details

`tatami.multiply` may not correctly propagate non-finite values when one of the matrices is sparse.

Value

For `tatami.dim`, an integer vector containing the dimensions of the matrix.

For `tatami.is.sparse`, a logical scalar indicating whether the matrix is sparse.

For `tatami.prefer.rows`, a logical scalar indicating whether the matrix prefers iteration by row.

For `tatami.get`, a numeric vector containing the contents of row or column `i`, respectively.

For `tatami.realize`, a numeric matrix or `dgCMatrix` with the matrix contents. The exact class depends on whether `x` refers to a sparse matrix.

For `tatami.multiply`, a numeric matrix containing the matrix product of `x` and `other`.

For `tatami.sums`, a numeric vector containing the row or column sums depending on `row=`.

For `tatami.variances`, a list containing the numeric vectors mean and variance. These containing the row/column means and variances, respectively, depending on `row=`.

For `tatami.sums.by.group`, a numeric matrix is returned.

- If `row=TRUE`, its number of rows is equal to that of `x` and its number of columns is equal to the number of groups. Each column corresponds to a group and contains the row sums across all columns of `x` assigned to that group.
- If `row=FALSE`, its number of columns is equal to that of `x` and its number of rows is equal to the number of groups. Each row corresponds to a group and contains the column sums across all columns of `x` assigned to that group.

For `tatami.medians`, a numeric vector containing the row or column medians, respectively.

For `tatami.nan.counts`, a numeric vector containing the number of NaNs in each row or column, respectively.

For all other functions, a new pointer to a matrix with the requested operations applied to `x` or `xs`.

Author(s)

Aaron Lun

Examples

```
x <- Matrix::rsparsematrix(1000, 100, 0.1)
ptr <- initializeCpp(x)
tatami.dim(ptr)
tatami.get(ptr, 1, row=TRUE)

rounded <- tatami.round(ptr)
tatami.get(rounded, 1, row=TRUE)

tatami.sums(ptr, row=FALSE, num.threads=2)
tatami.medians(ptr, row=FALSE, num.threads=2)
```

`toCsparse`*Convert a `SparseMatrix` to a `CsparseMatrix`*

Description

Exactly what it says in the title.

Usage

```
toCsparse(x)
```

Arguments

`x` Any object produced by block processing with `colBlockApply` or `rowBlockApply`. This can be a matrix, sparse matrix or a `SparseMatrix` object.

Details

This is intended for use inside functions to be passed to `colBlockApply` or `rowBlockApply`. The idea is to pre-process blocks for user-defined functions that don't know how to deal with `SparseMatrix` objects, which is often the case for R-defined functions that do not benefit from **beachmat**'s C++ abstraction.

Value

`x` is returned unless it is a `SparseMatrix` object, in which case an appropriate `CsparseMatrix` object is returned instead.

Author(s)

Aaron Lun

Examples

```
library(SparseArray)
out <- COO_SparseArray(c(10, 10),
  nzcoo=cbind(1:10, sample(10)),
  nzdata=runif(10))
toCsparse(out)
```

whichNonZero	<i>Find non-zero entries of a matrix</i>
--------------	--

Description

This function is soft-deprecated; users are advised to use [nzwhich](#) and [nzvals](#) instead.

Usage

```
whichNonZero(x, ...)
```

Arguments

x	A numeric matrix-like object, usually sparse in content if not in representation.
...	Further arguments, ignored.

Value

A list containing i, an integer vector of the row indices of all non-zero entries; j, an integer vector of the column indices of all non-zero entries; and x, a (usually atomic) vector of the values of the non-zero entries.

Author(s)

Aaron Lun

See Also

[which](#), obviously.

Examples

```
x <- Matrix::rsparsematrix(1e6, 1e6, 0.000001)
out <- whichNonZero(x)
str(out)
```

Index

apply, [4](#)
Arith, [10](#)
ArrayGrid, [4](#)

blockApply, [4](#), [5](#)

checkMemoryCache, [2](#)
colBlockApply, [3](#), [12](#)
Compare, [10](#)
currentViewport, [4](#)

DelayedMatrix, [4](#), [8](#)

flushMemoryCache (checkMemoryCache), [2](#)

getAutoBlockSize, [5](#)
getExecutor, [6](#)

initializeCpp, [2](#), [6](#), [8–10](#)
initializeCpp, ANY-method
 (initializeCpp), [6](#)
initializeCpp, ConstantArraySeed-method
 (initializeCpp), [6](#)
initializeCpp, DelayedAbind-method
 (initializeCpp), [6](#)
initializeCpp, DelayedAperm-method
 (initializeCpp), [6](#)
initializeCpp, DelayedMatrix-method
 (initializeCpp), [6](#)
initializeCpp, DelayedNaryIsoOp-method
 (initializeCpp), [6](#)
initializeCpp, DelayedSetDimnames-method
 (initializeCpp), [6](#)
initializeCpp, DelayedSubset-method
 (initializeCpp), [6](#)
initializeCpp, DelayedUnaryIsoOpStack-method
 (initializeCpp), [6](#)
initializeCpp, DelayedUnaryIsoOpWithArgs-method
 (initializeCpp), [6](#)
initializeCpp, dgCMatrix-method
 (initializeCpp), [6](#)

initializeCpp, dgeMatrix-method
 (initializeCpp), [6](#)
initializeCpp, dgRMatrix-method
 (initializeCpp), [6](#)
initializeCpp, externalptr-method
 (initializeCpp), [6](#)
initializeCpp, lgCMatrix-method
 (initializeCpp), [6](#)
initializeCpp, lgeMatrix-method
 (initializeCpp), [6](#)
initializeCpp, lgRMatrix-method
 (initializeCpp), [6](#)
initializeCpp, matrix-method
 (initializeCpp), [6](#)
initializeCpp, SVT_SparseMatrix-method
 (initializeCpp), [6](#)
is_sparse, [8](#)
isFileBackedMatrix
 (realizeFileBackedMatrix), [7](#)

Logic, [10](#)

Math, [10](#)

nzvals, [13](#)
nzwhich, [13](#)

path, [8](#)

realizeFileBackedMatrix, [7](#)
rowBlockApply, [12](#)
rowBlockApply (colBlockApply), [3](#)

SparseMatrix, [4](#), [5](#), [12](#)

tatami-utils, [8](#)
 tatami.arith (tatami-utils), [8](#)
 tatami.binary (tatami-utils), [8](#)
 tatami.bind (tatami-utils), [8](#)
 tatami.column (tatami-utils), [8](#)
 tatami.compare (tatami-utils), [8](#)

`tatami.dim (tatami-utils)`, 8
`tatami.get (tatami-utils)`, 8
`tatami.is.sparse (tatami-utils)`, 8
`tatami.log (tatami-utils)`, 8
`tatami.logic (tatami-utils)`, 8
`tatami.math (tatami-utils)`, 8
`tatami.medians (tatami-utils)`, 8
`tatami.multiply (tatami-utils)`, 8
`tatami.nan.counts (tatami-utils)`, 8
`tatami.not (tatami-utils)`, 8
`tatami.prefer.rows (tatami-utils)`, 8
`tatami.realize (tatami-utils)`, 8
`tatami.round (tatami-utils)`, 8
`tatami.row (tatami-utils)`, 8
`tatami.subset (tatami-utils)`, 8
`tatami.sums (tatami-utils)`, 8
`tatami.transpose (tatami-utils)`, 8
`tatami.variances (tatami-utils)`, 8
`toCsparse`, 5, 12

`which`, 13
`whichNonZero`, 13